Playing the Wrong Game: Smoothness Bounds for Congestion Games with Behavioral Biases

RESHEF MEIR, Technion—Israel Institute of Technology
DAVID PARKES, Harvard University

arXiv:1411.1751v4 [cs.GT] 21 May 2016

In many situations a player may act so as to maximize a perceived utility that is not exactly her utility function, but rather some other, biased, utility. Examples of such biased utility functions are common in behavioral economics, and include risk attitudes, altruism, present-bias and so on. When analyzing a game, one may ask how inefficiency, measured by the Price of Anarchy (PoA) is affected by the perceived utilities.

The smoothness method [Roughgarden and Tardos 2004; Roughgarden 2009] naturally extends to games with such perceived utilities or costs, regardless of the game or the behavioral bias. We show that such biased-smoothness is broadly applicable in the context of nonatomic congestion games. First, we show that on series-parallel networks we can use smoothness to yield PoA bounds even for diverse populations with different biases. Second, we identify various classes of cost functions and biases that are smooth, thereby substantially improving some recent results from the literature.

Additional Key Words and Phrases: Smoothness; Price of Anarchy; Behavioral Game Theory; Congestion Games

ACM Reference Format:
Reshef Meir and David Parkes, 2015. Playing the Wrong Game: Smoothness Bounds for Congestion Games with Behavioral Biases. *ACM Trans. Econ. Comp.* V, N, Article A (January YYYY), 31 pages.
DOI: 0000001.0000001

1. INTRODUCTION

Game theory is founded on the assumption that players are rational decision makers, i.e. maximizing their utility, and that groups of agents reach an equilibrium outcome. However agents, either human or automated, often have bounded resources that prevent them from finding the optimal response in every situation. Further, human decision makers are prone to various cognitive and behavioral biases, such as risk-aversion, loss-aversion, tendency to focus on short-term utility and so on.

As a concrete example, commuters may have some information on the expected congestion at each route via traffic reports or a cellphone app. However they also know that this information is inaccurate, and a risk-averse driver might take into account not just the expected congestion, but also the variance, standard deviation, a "safety margin" and so on. Further, different commuters may have different levels of risk-aversion.

The implications of these limitations and biases on game playing can be similarly described: from the perspective of an outside observer (who cares about the actual total latency), the players are playing the "wrong game," either by applying some simple heuristics, or by optimizing a different utility function from the one in the game specification [Devetag and Warglien 2008]. Moreover, different agents may have different perspective on the game, e.g. due to different levels of risk aversion.

Fortunately, as long as we can incorporate the biases that affect agents into suitably modified utility (or cost) functions, we can still use traditional equilibrium concepts. More formally, suppose that in the real underlying game G each agent i has some

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
© YYYY Copyright held by the owner/author(s). 1946-6227/YYYY/01-ARTA $15.00
DOI: 0000001.0000001

utility function u_i. Now, each player i sees her utility as some other function $\hat{u}_i$, and thus we are interested in the equilibria of the *biased game* $\hat{G}$ comprised of modified utilities $\{\hat{u}_i\}$. Crucially, we still measure the quality of the outcome on the real game G (see Footnote 3 for further discussion).

Congestion games are a good testbed for such ideas. First, they have very convenient theoretical properties, such as the existence of a potential function and the uniqueness of equilibrium (up to identical utilities, in nonatomic games). Second, equilibrium inefficiency and in particular the Price of anarchy is very well understood. Third, several cognitive and behavioral biases have already been suggested and studied in the context of congestion games (see Related Work).

Our starting point in this paper is the smoothness framework in nonatomic congestion games in the Wardrop [1952] model. The smoothness framework connects a particular property of the edge cost functions (say, $(1, \frac{1}{4})$-smoothness of affine functions), with a tight upper bound on the price of anarchy (PoA), this bound also being independent of the network topology [Roughgarden 2003].[1]

Smoothness for modified costs has also been considered; e.g. by Bonifaci et al. [2011] in the context of tolls that can also be viewed as a perturbation on utilities. The definition of smoothness naturally extends to account for bias, and the modified definition (which contains both the real cost c and the biased cost $\hat{c}$) guarantees a "biased" price of anarchy (BPoA) bound in a similar manner to standard smoothness PoA analysis. Moreover, it is easy to see that while [Bonifaci *et al.* 2011] considered a particular modified cost function, the same biased-smoothness approach works for any combination of c and $\hat{c}$. BPoA bounds have been similarly attained via smoothness for finite games with altruistic players [Chen *et al.* 2011].

Thus the PoA is not only robust to the network topology and to different notions of equilibrium [Roughgarden 2009], but also to players that are not playing exactly by the game specification. This observation brings about two important challenges: First, it is unlikely that all of the participants in a game have exactly the same bias (see example on risk-aversion). Thus we want to combine smoothness results of several different types to attain a single (biased) PoA bound. Second, can we find broad classes of games and behavioral biases for which smoothness applies?

1.1. Paper structure and Contribution

In Section 2 we briefly go over the standard model of nonatomic congestion games, price of anarchy and smoothness. In Section 3 we extend the definition of a smooth cost function to a more general notion of biased-smoothness, following similar extensions for specific modified costs [Bonifaci *et al.* 2011; Chen *et al.* 2011]. We then provide positive results in both of the above directions:

- For games over series-parallel networks, players are affected almost only by their own bias, as the negative externalities due to the biases of other agent types are bounded.
- We conclude that under the same structural restriction, the BPoA in a game with a mixed population of different behavior types is bounded in terms of the biased smoothness parameters of all the participating types. We emphasize that this holds regardless of the specific biases, as long as for each agent type we have established smoothness (Section 4).

[1] While historically the results of Roughgarden [2003] preceded the introduction of smoothness in games, today the all of these results are commonly interpreted under the smoothness framework, see Roughgarden [2009].

- In contrast, without such a structural restriction, even a small fraction of "bad" agents may inflict unbounded damage on the society.
- We prove biased-smoothness for several classes of cost functions in non-atomic congestion games with particular behavioral biases:
  - We derive tight BPoA bounds for agents with tax sensitivity under polynomial cost functions of any degree $d \geq 1$ (Section 5).
  - We improve the upper bounds of Meir and Parkes [2015] for pessimist agents (Section 6).
  - We significantly improve the upper bound of Nikolova and Stier-Moses [2015] for risk-averse agents, showing it is independent of the graph topology (Section 7).

We conclude with a brief comparison to related work, a discussion of our results, and suggest some directions for future work.

2. PRELIMINARIES

Nonatomic routing games. Following the definitions of Roughgarden [2003] and Roughgarden and Tardos [2004], a *nonatomic routing game* (NRG) is a tuple $G = \langle V, E, m, \mathbf{c}, \mathbf{u}, \mathbf{v}, \mathbf{n} \rangle$, where

- (V, E) is a directed graph;
- $m \in \mathbb{N}$ is the number of agent types;
- $\mathbf{c} = (c_e)_{e \in E}$, $c_e(x) \geq 0$ is the cost incurred when x agents use edge e;[2]
- $\mathbf{u}, \mathbf{v} \in V^m$, where (u_i, v_i) are the source and target nodes of type i agents;
- $\mathbf{n} \in \mathbb{R}_+^m$, where $n_i \in \mathbb{R}_+$ is the total mass of type i agents. The total mass of agents of all types is $\sum_{i \leq m} n_i = n$, and we assume unless specified otherwise that $n = 1$.

We denote by $A_i \subseteq 2^E$ the set of all directed paths between the pair of nodes (u_i, v_i) in the graph. Thus A_i is the the set of *actions* available to agents of type i. We denote by $A = \cup_i A_i$ the set of all directed source-target paths. We assume that the costs c_e are non-decreasing, continuous, differentiable and semi-convex (i.e., $xc_e(x)$ is convex). Such cost functions are called *standard* [Roughgarden 2003].

Player-specific costs. A *nonatomic routing game with player-specific costs* (PNRG) is a tuple $G = \langle V, E, m, (\mathbf{c}^i)_{i \leq m}, \mathbf{u}, \mathbf{v}, \mathbf{n} \rangle$. The difference from a NRG is that agents of each type i experience a cost of $c_e^i(x)$ when x agents use edge e. We can have multiple types with the same source and target nodes to allow diversity of behavior. To avoid confusion, we refer to (v_i, u_i) (or A_i) as the *demand type* and to $\mathbf{c}^i$ as the *cost type*. Thus the type i specifies both the demand type and the cost type.

A PNRG is *symmetric* if all agents have the same demand type, i.e., $A_i = A$ for all i. A PNRG is a *resource selection game* (RSG) if (V, E) is a graph of parallel links. That is, if the action of every agent is to select a single edge.

Game states. A *state* (or action profile) of a PNRG is a vector $\mathbf{s} \in \mathbb{R}_+^{|A| \times m}$, where $s_{f,i}$ is the amount of agents of type i that use path $f \in A_i$. In a valid state, $\sum_{f \in A_i} s_{f,i} = n_i$ for all i. The total traffic on path $f \in A$ is denoted by $s_f = \sum_{i=1}^m s_{f,i}$. Similarly, the total traffic on edge $e \in E$ is denoted by $s_e = \sum_{f : e \in f} s_f$.

The *social cost* in a profile $\mathbf{s}$ in a NRG G is defined as $SC(G, \mathbf{s}) = \sum_{i=1}^m \sum_{f \in A_i} s_{f,i} c(f, \mathbf{s})$. An equivalent way to write the social cost is as

$$SC(G, \mathbf{s}) = \sum_{e \in E} s_e c_e(s_e) \tag{1}$$

[2] c_e are sometimes referred to as *latency functions*.

This shows that the social cost only depends on the total traffic per edge.

Remark 2.1. This social cost provides an *objective measure* of the total latency in the system. However, the players in our model may act based on different cost functions, either because they have a different objective than simply minimizing their latency, or because they are affected by cognitive and behavioral biases. We provide examples of both types of discrepancies in Section 3.1.

We omit the argument G when it is clear from the context. We denote by $OPT(G)$ some profile with minimal total cost, i.e. $OPT(G) \in \operatorname{argmin}_{\mathbf{s}} SC(G, \mathbf{s})$.

Equilibrium. A state $\mathbf{s}$ for an NRG is an *equilibrium* in game G if for every agent type i and actions $f_1, f_2 \in A_i$ with $s_{i,f_1} > 0$, we have $c(f_1, \mathbf{s}) \leq c(f_2, \mathbf{s})$. That is, if no agent can switch to a path with a lower cost. This provides the analogy of a Nash equilibrium for nonatomic games, and it extends naturally for player-specific NRGs.

It is known that in any NRG there is at least one equilibrium, and that this can be reached by a simple best-response dynamic. Further, all equilibria have the same social cost and in every equilibrium all agents of type i experience the same cost [Aashtiani and Magnanti 1981; Milchtaich 2000; Roughgarden and Tardos 2004]. Player-specific NRGs are also guaranteed to have at least one equilibrium [Schmeidler 1973], however, equilibrium costs may not be unique, and best-response dynamics may not converge.

Affine routing games. In an *affine* NRG, all cost functions take the form of a linear function. That is, $c_e(x) = a_e x + b_e$ for some constants $a_e \geq 0, b_e \geq 0$. The social cost can be written as $SC(G, \mathbf{s}) = \sum_{e \in E} a_e (s_e)^2 + b_e s_e$. ***Pigou's example*** is the special case of an affine RSG with two resources, where $c_1(x) = 1$ and $c_2(x) = ax$. We denote by $G_P(a, 1)$ the instance where $c_2(x) = ax$ (and $c_1(x) = 1$), and $G_P(a, d)$ the instance where $c_2(x) = ax^d$ (and $c_1(x) = 1$).

The price of anarchy. Let $EQ(G)$ be the set of equilibria in game G. The *price of anarchy* (PoA) of a game is the ratio between the social cost in the worst equilibrium in $EQ(G)$ and the optimal social cost [Koutsoupias and Papadimitriou 1999]. Since all equilibria have the same cost, we can write $\text{PoA}(G) = \frac{SC(G, \mathbf{s}^*)}{SC(G, OPT(G))}$, where $\mathbf{s}^*$ is any equilibrium of G. For example in affine NRGs, it is known that $\text{PoA}(G) \leq \frac{4}{3}$, and this bound is attained by the Pigou example $G_P(1, 1)$ [Roughgarden and Tardos 2004].

Smoothness. A cost function c is (λ, μ)-smooth for $\lambda \geq 0, \mu < 1$ if for any $x, x' \geq 0$, it holds that

$$c(x)x' \leq \lambda x' c(x') + \mu x c(x). \tag{2}$$

A game G is (λ, μ)-smooth if all cost functions in G are (λ, μ)-smooth. For any (λ, μ)-smooth game, $\text{PoA}(G) \leq \frac{\lambda}{1-\mu}$ [Correa *et al.* 2008; Roughgarden 2009]. Moreover, w.l.o.g. $\lambda = 1$ (that is, for any class of cost functions there is an optimal pair $(1, \mu)$ for some μ).

For example, affine routing games are $(1, \frac{1}{4})$-smooth [Roughgarden 2009], providing an alternative derivation of the PoA bound of $\frac{4}{3}$.

3. BIASED COSTS

Given an NRG $G = \langle V, E, m, \mathbf{c}, \mathbf{u}, \mathbf{v}, \mathbf{n} \rangle$ and modified cost functions $\hat{c}_e^i$ for every type $i \leq m$ and edge $e \in E$, we obtain a PNRG $\hat{G} = \left\langle V, E, m, (\mathbf{c}^i)_{i \leq m}, \mathbf{u}, \mathbf{v}, \mathbf{n} \right\rangle$. We assume that players act based on their modified cost functions, irrespective of whether this is a rational behavior or not (see Remark 2.1). We refer to $\hat{G}$ as the *biased game*, where every agent of type i experiences a cost of $\hat{c}_e^i$ rather than c_e .

Both games G and $\hat{G}$ have a role in our model.

Definition 3.1. A *nonatomic routing game with biased costs* is a tuple $\mathcal{G} = \left\langle G, \hat{G} \right\rangle$, where $\hat{G}$ is a biased version of G.

The way we interpret $\mathcal{G}$ is that players play the wrong game $\hat{G}$ (and thus, it is the equilibria of this game that matter), whereas their true costs are according to game G. In other words, the biased costs $\hat{c}_e^i$ affect behavior while the actual cost in state s is the original true cost $c(\mathbf{s})$, and it is this true cost that determines welfare).[3]

Biased Price of Anarchy. We measure the price of anarchy in a game with biased costs by comparing the equilibria of $\hat{G}$ to the optimum of G. Formally:

$$\text{BPoA}(G, \hat{G}) = \sup_{\mathbf{s} \in EQ(\hat{G})} \frac{SC(G, \mathbf{s})}{SC(G, OPT(G))}. \tag{3}$$

In the uniform bias case where $\hat{\mathbf{c}}^i = \hat{\mathbf{c}}$ for all i, the game $\hat{G}$ is just another NRG, and thus has a potential function.

In the general case, an equilibrium of $\hat{G}$, and thus of $\mathcal{G}$, always exists, although $\hat{G}$ may not be a potential game. Existence follows from general existence results on nonatomic games with convex strategy spaces and continuous utilities [Schmeidler 1973].

3.1. Biased Cost Functions

Uncertainty and Risk Aversion. A simple form of biased cost is induced by what we term *pessimism*, and can be thought of as a special case of risk-aversion. Suppose that whenever faced with some load s_e on edge e, an agent believes this number to be inaccurate, and takes a safety margin by playing as if the actual load is $r \cdot s_e$ (for some fixed private parameter $r \geq 1$).

Such an agent will play *as if* every cost function c_e is replaced with a new cost function $\hat{c}_e^r$, where $\hat{c}_e^r(x) = c_e(rx)$ (see Fig. 1(b)). This is equivalent to the worst-case-cost (WCC) model of Meir and Parkes [2015], which models agents with strict uncertainty over the actual costs.

A somewhat different model considers the uncertainty induced by an arbitrary noise distribution $\epsilon_e(x)$ (with mean 0), by setting $\hat{c}_f^\gamma(\mathbf{s}) = c_f(\mathbf{s}) + \gamma \text{var}(\sum_{e \in f} \epsilon_e(s_e))$, indicating that agents prefer paths with a lower variance [Nikolova and Stier-Moses 2015]. Since the variance is additive along the path, we can rewrite the edge cost as $\hat{c}_e^\gamma(x) = c_e(x) + \gamma \text{var}(\epsilon_e(x))$. We assume that $\text{var}_e(x)$ is bounded by $\kappa c_e(x)$ for some constant $\kappa > 0$ (see Fig. 1(d) for an example that violates this assumption).

The previous examples of biased games assumed players' perceptions that are biased or misguided. The next two examples are cases where the objectives of the players and the analyst/designer differ (see Remark 2.1).

Sensitivity to tax. Suppose that a central authority imposes an optimal tax of $xc_e'(x)$ on every edge (where $c_e'(x)$ is the derivative of c_e). If agents treat the monetary tax

[3] One may argue that since agents are acting so as to minimize c^i rather than c, this is the cost we should take into account when computing welfare. We justify using G as the baseline for welfare in one of two ways: (a) under a bounded-rationality approach (e.g., assuming that agents are unaware of the true cost, follow a misguided heuristic, or assigning excessive probabilities to rare outcomes); (b) when the social cost is the goal used by the system designer or analyst, e.g. to minimize the average latency regardless of how the agents feel about the outcome (see also Remark 2.1).

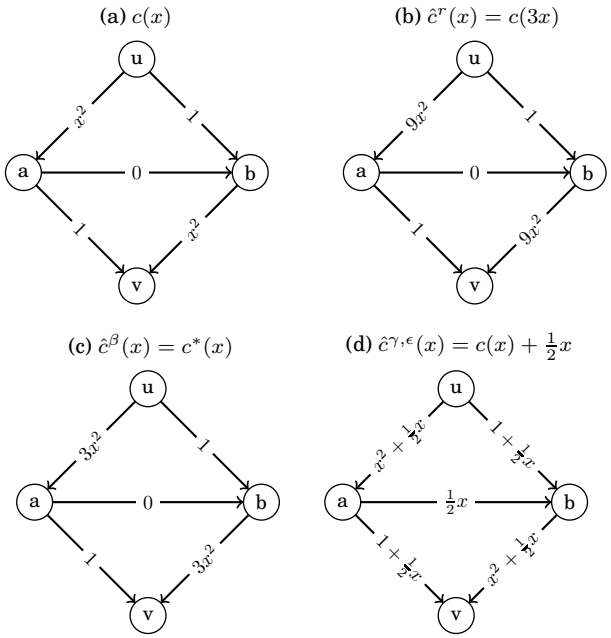

Figure 1: Fig. 1a illustrates a game G that is a quadratic variation of the Braess paradox [Frank 1981]. The other subfigures present three biased versions of G. For example, the game $\hat{G}_b$ is played by pessimistic agents, with parameter $r = 3$. Hence the edges with fixed costs do not change, but the biased cost on the edge a-b for example is $\hat{c}_{u-a}(x) = c_{u-a}(3x) = (3x)^2 = 9x^2$.

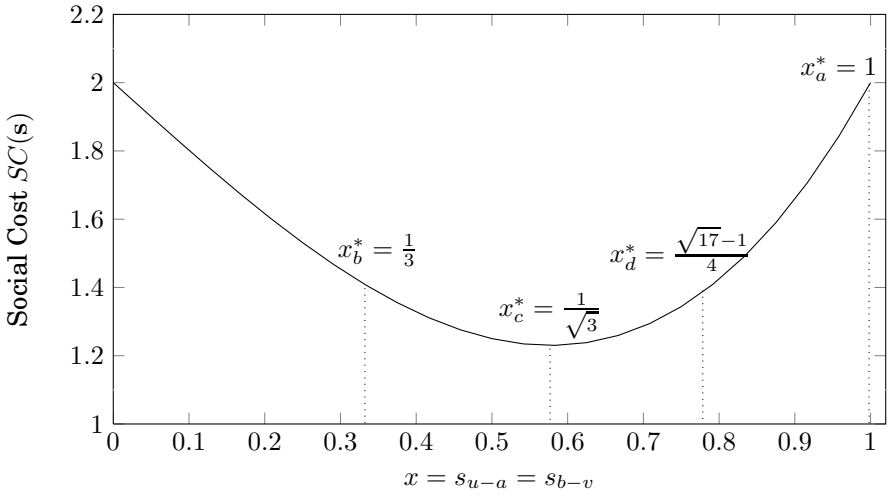

Figure 2: The social cost in game G from Fig. 1 as a function of the traffic on the edge $u - a$ (assuming symmetric flow). The equilibria of all four games are presented.

and the cost due to delay in the same way, then we get biased cost function $\hat{c}_e(x) = c_e(x) + xc'_e(x) = c^*_e(x)$, with the effect that the incentives are perfectly aligned with those of the society, and the PoA is 1 [Beckmann *et al.* 1956]. We emphasize that the designer in this case is only interested in the cost stemming from congestion, and disregards any monetary transfers.

However since the basic cost is in terms of delay, and the tax is in terms of money, different agents may have different money value for time, and hence different tax sensitivity [Cole *et al.* 2003; Karakostas and Kolliopoulos 2004; Fleischer *et al.* 2004; Fotakis *et al.* 2010]. This suggests a biased cost function $\hat{c}^\beta(x) = c(x) + \beta xc'(x)$ (see Fig. 1(c)). As discussed in [Yang and Zhang 2008] these individual differences may be *unobservable*, and hence cannot be taken into account by the tax mechanism.

In our case, the designer does not know the value of β (or its distribution), and does not try to fit the tolls to the present population. Rather, we are interested in the effect of the standard, marginal-cost toll scheme (i.e., optimal for unbiased agents) on the social cost as β varies.

Capacitated resources. Suppose that each edge has a hard capacity limit L_e, meaning that a state s is valid only if $s_e \leq L_e$ for all $e \in E$. We can translate this restriction to an endogenous bias, where $\hat{c}_e(x) = c_e(x)$ if $x \leq L_e$, and some large constant $M > 0$ otherwise. Since this would make $\hat{c}_e$ discontinuous, we can define a similar function that only starts to increase after $L_e - \delta$, for example $\hat{c}^\delta_e(x) = c_e(x)(yM + 1)y^2$ for all $x \geq L_e\delta$, where $y = \frac{x-(L_e-\delta)}{\delta}$. This means that c^δ_e is standard, equals to c_e in the range $x \leq L_e$, and is greater than M in the range $x \geq L_e$. As $\delta \to 0$, game $\hat{G}$ approaches the capacitated game. The BPoA then compares the cost of an equilibrium subject to capacity constraints to that of the optimal uncapacitated flow.[4]

Other types of biases that can be similarly modeled but are not studied in our paper are idiosyncratic preferences to certain strategies [Sandholm 2007], and altruistic behavior [Caragiannis *et al.* 2010; Chen *et al.* 2011].

3.2. Smoothness for Biased Costs

Our goal is to provide bounds on the biased Price of Anarchy for a given game with biased costs $\left\langle G, \hat{G} \right\rangle$. That each of c_e and $\hat{c}^i_e$ are smooth is insufficient to provide such a bound.

Example 3.2. Consider a game with pessimistic agents $\mathcal{G}_r = \left\langle G_r, \hat{G}^r_r \right\rangle$ where $G_r = G_P(2/r, 1)$, i.e., the Pigou example. For any r the equilibrium of $\hat{G}^r_r$ (and thus of $\mathcal{G}_r$) is the same: 1/2 of all agents use each resource. However as r increases, the optimal solution of G shifts more agents to the second resource, and the optimal social cost decreases to 0. Thus the gap between the equilibrium cost and the optimal cost (the BPoA) goes to infinity with r.

We extend the definition of smoothness to games with biased costs in a way that takes into account both c and $\hat{c}$. This technique has been applied before for specific modified costs, for example nonatomic games with restricted taxes [Bonifaci *et al.* 2011] and atomic games with altruistic players [Chen *et al.* 2011]. We provide an extension that handles general biases.

[4]Efficiency loss with capacitated resources is considered in Bonifaci et al. [2011], and is not within the scope of this paper, except for showing that it fits under the framework of biased PoA.

Definition 3.3. The function c is (λ, μ)-*biased-smooth* w.r.t. biased cost function $\hat{c}$, if for any $x, x' \in \mathbb{R}_+$,

$$c(x)x + \hat{c}(x)(x' - x) \leq \lambda c(x')x' + \mu c(x)x. \tag{4}$$

It is instructive to check the familiar case where there is no bias. Indeed, if $\hat{c} = c$, and c is (λ, μ)-smooth, then

$$c(x)x + \hat{c}(x)(x' - x) = c(x)x + c(x)(x' - x) = c(x)x' \leq \lambda c(x')x' + \mu c(x)x,$$

and Eq. (4) collapses to the standard definition of smoothness from Eq. (2).

Recall that the PoA of a (λ, μ)-smooth game is bounded by $\frac{\lambda}{1-\mu}$. This bound extends to games with biased costs that are biased-smooth and when all agents have the same bias.

THEOREM 3.4 ([CORREA *et al.* 2008; ROUGHGARDEN 2009; BONIFACI *et al.* 2011]). *Consider a game with biased costs $\mathcal{G}$ where every cost function c_e is $(\hat{\lambda}, \hat{\mu})$-biased smooth w.r.t. biased cost function $\hat{c}^i_e$. Let $\mathbf{s}^*$ be any equilibrium of the game $\hat{G}$, and $\mathbf{s}'$ any valid state. Then $SC(G, \mathbf{s}^*) \leq \frac{\hat{\lambda}}{1-\hat{\mu}} SC(G, \mathbf{s}')$.*

PROOF. The proof extends the standard proof of PoA bounds for nonatomic congestion games via smoothness arguments. In any equilibrium $\mathbf{s}^*$, the *variational inequality* $\sum_e \hat{c}_e(s^*_e)s^*_e \leq \sum_e \hat{c}_e(s^*_e)s'_e$ holds (see [Correa *et al.* 2008; Roughgarden 2009]). Thus,

$$\begin{aligned} SC(G, \mathbf{s}^*) &\leq SC(G, \mathbf{s}^*) + \sum_e \hat{c}_e(s^*_e)s'_e - \sum_e \hat{c}_e(s^*_e)s^*_e \\ &= \sum_e [c_e(s^*_e)s^*_e + \hat{c}_e(s^*_e)(s'_e - s^*_e)] \\ &\leq \sum_e [\hat{\lambda} c_e(s'_e)s'_e + \hat{\mu} c_e(s^*_e)s^*_e] = \hat{\lambda} SC(G, \mathbf{s}') + \hat{\mu} SC(G, \mathbf{s}^*). \qquad \text{(by Def. 3.3)} \end{aligned}$$

We get the bound by reorganizing the terms. The only part that differs from the standard smoothness is the first inequality. □

COROLLARY 3.5. *Consider a game G where each cost function c_e is $(\hat{\lambda}, \hat{\mu})$-biased smooth w.r.t. $\hat{c}_e$. Then BPoA$(G, \hat{G}) \leq \frac{\hat{\lambda}}{1-\hat{\mu}}$.*

The equilibrium of each of the biased games from Fig. 1 are presented in Fig. 2. The equilibrium of the real game $\mathbf{s}^*_a$ is suboptimal, as all agents take the long path $u - a - b - v$. Thus BPoA$(G, \hat{G}_a) =$ PoA$(G) \cong \frac{2}{1.23} = 1.625$. Similarly, in the equilibrium $\mathbf{s}^*_b, \mathbf{s}^*_d$ there is too little or too much traffic on the long path, respectively.

An alternative derivation of optimal taxation. We can also check that the extension provides the right result in regard to modified costs $c^*(x) = c(x) + c'(x)x$ that represent optimal taxation and should lead to optimal play [Beckmann *et al.* 1956]. Let's confirm this result via a biased-smoothness argument.

LEMMA 3.6. *If $c(x)x$ is convex and we set modified cost $\hat{c}(x) = c^*(x) = c(x) + c'(x)x$, then cost function c is $(1, 0)$-biased smooth w.r.t. to $\hat{c}$.*

PROOF. Let $h(x) = xc(x)$. By convexity of h,

$$c(x)x + \hat{c}(x)(x' - x) = h(x) + h'(x)(x' - x) \leq h(x') = x'c(x'),$$

thereby affirming that the price of anarchy of $\langle G, G^* \rangle$ is 1. □

For example, the biased game $\hat{G}_c$ in Fig. 1 has optimal taxation, and we can see in Fig. 2 that its equilibrium x_c^* minimizes the social cost.

Robustness to small biases. Consider some bias s.t. $\max\{\frac{\hat{c}(x)}{c(x)}, \frac{c(x)}{\hat{c}(x)}\} \leq 1+\epsilon$ for all $x \geq 0$. Then

$$c(x)x + \hat{c}(x)(x'-x) \leq c(x)x + c(x)(x'-x) + |\epsilon c(x)(x'-x)| \leq (1+\epsilon)(c(x)x + c(x)(x'-x)).$$

Thus if c is (λ,μ)-smooth, then c is $((1+\epsilon)\lambda, (1+\epsilon)\mu)$ biased-smooth w.r.t. $\hat{c}$, and the effect on PoA is negligible for small ϵ (provided that μ is not very close to 1).

4. DIVERSE POPULATION

In general, there may be many types of agents in a game, each with a different behavioral bias. For example we can have agents with various levels of tax-sensitivity or pessimism, a small number of adversarial agents, and so on. Our main goal is to provide claims about the welfare achieved in the equilibria that correspond to such populations.

We define the following variations of a game G, given a set of biased cost functions $(\mathbf{c}^i)_{i\leq m}$:

- G^i is obtained from G by setting the *demand type* of all agents to A_i. Game G^i is a symmetric NRG, where all costs are as in G. Also $G^i = G$ iff G is symmetric.
- $\hat{G}^i$ is obtained from G^i by setting the *cost type* of all players to $\hat{\mathbf{c}}^i$. Game $\hat{G}^i$ is a symmetric NRG. This corresponds to game with biased costs $\mathcal{G}^i = \left\langle G, \hat{G}^i \right\rangle$, where all agents have the same bias.

Also recall our previous definitions:

- $\hat{G}$ is a PNRG that is obtained from G by setting the cost type of each type i agent to $\hat{\mathbf{c}}^i$. This corresponds to game with biased costs $\mathcal{G} = \left\langle G, \hat{G} \right\rangle$, where the modified cost function is obtained for each agent type via a distinct, type-specific mapping.

Suppose we can show biased smoothness of $(\hat{\lambda}^i, \hat{\mu}^i)$ for each of the uniform games with biased costs $\mathcal{G}^i$. The primary question is whether we can get a bound on BPoA($\mathcal{G}$) in terms of the smoothness parameters, $(\hat{\lambda}^i, \hat{\mu}^i)_{i\leq m}$, for each type. Unfortunately, in the most general case of nonatomic congestion games even a small fraction of "bad" agents can arbitrarily increase the BPoA (see Section 4.4).

To obtain a positive result that remains general to different cost functions and behavioral biases, we restrict the structure of the network to series-parallel graphs. In Section 6, we show a bound for general networks, when we consider more restricted cost functions and biases.

4.1. Directed Series-Parallel Graphs

A *directed series-parallel (DSP) graph* is an acyclic directed graph (V, E, u, v) with a source u and a target v, and is composed recursively.

The basic graph is a single directed edge $u - v$. Two DSPs $(V_1, E_1, u_1, v_2), (V_2, E_2, u_2, v_2)$ can be composed to obtain a new DSP either in a *serial way*: merge u_1 with v_2; or in a *parallel way* by merging u_1 with u_2, and v_1 with v_2. A (non-directed) series-parallel graph is obtained by taking a DSP and removing edge orientations.

A *directed series-parallel game* (DSPG) is a nonatomic congestion game over a DSP graph. RSGs are a special case of DSPGs in which all edges are composed in parallel.

We assume for simplicity that all cost functions in DSPGs are either strictly increasing or constant. In addition, and w.l.o.g., we assume there are no two paths with the same constant cost between any pair of nodes (otherwise we can treat them as a single path).

Let $s_{i,e}$ be the mass of type i agents on path f in state $\mathbf{s}$. We have $n_i = \sum_{f \in A} s_{i,f}$ for any state $\mathbf{s}$ and type i. Denote the *support of type i strategy* by $S_i = \{f \in A : s_{i,f} > 0\}$. In equilibrium $\mathbf{s}^*$, for all $f \in S_i$, $\hat{c}^i_f(\mathbf{s}^*) = \sum_{e \in f} \hat{c}^i_e(s^*_e)$ is the same, and weakly lower than the biased cost of any path $f' \notin S_i$. We denote by $SC_i(G, \mathbf{s}) = \sum_{f \in S_i} s_{i,f} c_f(\mathbf{s})$ the cost experienced by type i agents in state $\mathbf{s}$. By summing over all types, we get the social cost:

$$\sum_{i \le m} SC_i(G, \mathbf{s}) = \sum_{i \le m} \sum_{f \in S_i} s_{i,f} c_f(\mathbf{s}) = \sum_{i \le m} \sum_{f} s_{i,f} c_f(\mathbf{s})$$
$$= \sum_{f} \sum_{i \le m} s_{i,f} c_f(\mathbf{s}) = \sum_{f} s_f c_f(s_f) = SC(G, \mathbf{s})$$

Our proof strategy is to first show a useful property of an equilibrium in series-parallel graphs. Then we will proceed to bound the cost for each type ($SC_i(G, \mathbf{s})$) separately and combine the bounds.

Let $\mathbf{s}^*$ be an equilibrium of a DSPG with biased costs $\hat{G}$, and $\mathbf{s}^i$ be an equilibrium of the game $\hat{G}^i$ (i.e., the game in which the demand and cost types of all players are those of type i agents).

LEMMA 4.1. *For any used path $f \in S^*_i$, $s^i_f \ge s^*_f$. In particular, for any $e \in S^*_i$, $s^i_e \ge s^*_e$.*

PROOF. The proof is by induction on the number of merging steps in the construction of the underlying graph. The base case for an edge is immediate.

If the DSP is composed of two smaller DSPs in series, then we can simply analyze each one as a separate game. An equilibrium of the full game projected over each subgame does not change, and thus we get the lemma.

For composition in parallel, the graph (V, E) is attained my merging two graphs (V_1, E_1) and (V_2, E_2) with the same source and target u, v. Denote the observed cost to type i agents in $\mathbf{s}^*$ (regardless of their chosen path) by $c^* = \hat{c}^i(f, \mathbf{s}^*)$.

By construction, there is no directed path in (V, E) that contains edges from both E_1, E_2. Denote by n^*_1, n^*_2 the populations of agents in $\mathbf{s}^*$ choosing paths in E_1, E_2, respectively (note that these are disjoint sets and $|n^*_1| + |n^*_2| = n$). Let G_1 be the game played on (V_1, E_1) by agents n^*_1 and likewise for (V_2, E_2).

Consider the equilibrium $\mathbf{s}^i$ (in the complete game) when all agents are of type i. It holds that $s^i_{f'} > 0$ for some $f' \in A_i$, or otherwise $\mathbf{s}^i$ is not an equilibrium. W.l.o.g. $f' \subseteq E_1$, meaning that n^*_1 contains some type i agents.

Similarly partition the population into n^i_1, n^i_2 according the their path choice in $\mathbf{s}^i$.[5] Assume toward a contradiction that $s^i_{f'} < s^*_{f'}$ for some $f' \in S^*_i$, where w.l.o.g. $f' \subseteq E_1$. Since the total mass of agents in fixed, some other $u_i - v_i$ paths must have more agents in $\mathbf{s}^i$. Crucially, $\mathbf{s}^i$ remains the equilibrium of each subgame, that is, if we consider the agents of n^i_1 playing on (V_1, E_1) and likewise for the other subgame.

If $|n^i_1| \ge |n^*_1|$, then consider the game G_1 alone with $|n^*_1|$ agents of type i. Denote by $\mathbf{s}^{(1)}$ the attained equilibrium, which we now compare with $\mathbf{s}^i$ (on the graph (V_1, E_1)).

[5]Note that if $v_i \neq v$ or $u_i \neq u$ then $n^i_1 = n$.

Since we only decreased the number of agents, and in both games they are all of type i, we get an equilibrium $s_f^{(1)} \leq s_f^i$ for all f. In particular, $s_{f'}^{(1)} \leq s_{f'}^i < s_{f'}^*$. However by our induction hypothesis applied to game G_1, we should have $s_f^{(1)} \geq s_f^*$ for any path $f \subseteq A_i \cap E_1$, and for f' in particular. This is a contradiction.

Thus $|n_1^i| < |n_1^*|$, which means $|n_2^i| > |n_2^*|$. In particular there is some path $f'' \subseteq E_2$ s.t. $s_{f''}^i > s_{f''}^* \geq 0$.

We know that all agents in $\mathbf{s}^i$ observe the same cost $c^i = \hat{c}^i(f', \mathbf{s}^i) = \hat{c}^i(f'', \mathbf{s}^i)$ (as $s_{f''}^i > 0, s_{f'}^i > 0$), whereas for any other path f (either in E_1 or E_2) we have $\hat{c}^i(f, \mathbf{s}^i) \geq c^i$. Also $\hat{c}^i(f', \mathbf{s}^*) = c^*$ and the cost in $\mathbf{s}^*$ is at least c^* for any path including f''. Therefore:

$$c^i = \hat{c}^i(f', \mathbf{s}^i) = \hat{c}^i(f'', \mathbf{s}^i) \geq \hat{c}^i(f'', \mathbf{s}^*) \geq c^* = \hat{c}^i(f', \mathbf{s}^*) \geq \hat{c}^i(f', \mathbf{s}^i) = c_i.$$

This entails that (a) the costs of all edges along the paths f', f'' are constants (as they are not strictly increasing); (b) that the paths f', f'' have the same constant cost. This is in contradiction to our earlier assumption. □

4.2. Fraction-independent bound

The first bound that we obtain is independent of the fractions of each type of agent in the population, and is useful when there is a large group of agents of the same type.

PROPOSITION 4.2. *Let $\mathbf{s}^*$ be an equilibrium of DSPG $\hat{G}$. For any type i, $SC_i(G, \mathbf{s}^*) \leq \frac{\hat{\lambda}^i}{1-\hat{\mu}^i} SC(G, OPT(G^i))$.*

PROOF. Recall that $\mathbf{s}^i$ is an equilibrium of $\hat{G}^i$.

$$\begin{aligned} SC_i(G, \mathbf{s}^*) &= \sum_{f \in A_i} s_{i,f}^* c(f, \mathbf{s}^*) = \sum_{f \in S_i^*} s_{i,f}^* c(f, \mathbf{s}^*) \\ &\leq \sum_{f \in S_i^*} s_f^i c(f, \mathbf{s}^i) && \text{(From Lemma 4.1)} \\ &\leq \sum_{f \in A_i} s_f^i c(f, \mathbf{s}^i) = SC(G, \mathbf{s}^i) = BPoA(G, \hat{G}^i) SC(G, OPT(G^i)) && (5) \\ &\leq \frac{\hat{\lambda}^i}{1-\hat{\mu}^i} SC(G, OPT(G^i)). \quad \square && \text{(From Cor. 3.5)} \end{aligned}$$

That is, if we only look at the cost to type i agents under the equilibrium $\mathbf{s}^*$ of the game with biased costs, it is not worse than the overall cost in the equilibrium $\mathbf{s}^i$ of the game $\hat{G}^i$ when all agents are of type i. In addition, the overall cost of this equilibrium is not much worse than the optimal social cost in this game $\hat{G}^i$ when all agents are of type i.

Prop. 4.2 has two immediate corollaries. First, the presence of a small fraction α of agents with arbitrary or adversarial behavior can harm the remaining agents by a factor of no more than $1/(1-\alpha)$:

THEOREM 4.3. *Suppose that a game with biased costs $\mathcal{G} = \left\langle G, \hat{G} \right\rangle$ is a symmetric game over a* DSP network*, where the total fraction of agents types other than i, $\sum_{j \neq i} n_j \leq \alpha$, for some $0 < \alpha < 1$. Let $\mathbf{s}^*$ be an equilibrium of $\mathcal{G}$. Then there is some optimal state $\tilde{\mathbf{s}} \in \operatorname{argmin}_{\mathbf{s}} SC(G, \mathbf{s})$ s.t. $SC_i(G, \mathbf{s}^*) \leq \frac{1}{1-\alpha} \frac{\hat{\lambda}^i}{1-\hat{\mu}^i} SC_i(G, \tilde{\mathbf{s}})$.*

As before, $\hat{\lambda}^i$ and $\hat{\mu}^i$ are two constants s.t. each c_e is $\frac{\hat{\lambda}^i}{1-\hat{\mu}^i}$ w.r.t. c_e^i.

PROOF. Let $\tilde{\mathbf{s}}$ be the optimal state where type i agents suffer the highest cots (since G is symmetric, we can swap agents arbitrarily without changing the social cost). Since G is symmetric, $G^i = G$. Then by Prop. 4.2,

$$SC_i(G, \mathbf{s}^*) \leq \frac{\hat{\lambda}^i}{1-\hat{\mu}^i} SC(G, OPT(G)) = \frac{\hat{\lambda}^i}{1-\hat{\mu}^i} SC(G, \tilde{\mathbf{s}})$$
$$\leq \frac{\hat{\lambda}^i}{1-\hat{\mu}^i} \frac{n}{n_i} SC_i(G, \tilde{\mathbf{s}}) \leq \frac{1}{1-\alpha} \frac{\hat{\lambda}^i}{1-\hat{\mu}^i} SC_i(G, \tilde{\mathbf{s}}),$$

as required. □

We also obtain a price of anarchy bound, by summing Eq. (5) over all types. This bound is useful when there are only a few types of agents.

THEOREM 4.4. *Suppose that* $\mathcal{G} = \left\langle G, \hat{G} \right\rangle$ *is a symmetric game over a* DSP network. $\mathit{BPoA}(\mathcal{G}) \leq \sum_{i \leq m} \mathit{BPoA}(G, \hat{G}^i) \leq \sum_{i \leq m} \frac{\hat{\lambda}^i}{1-\hat{\mu}^i}$.

4.3. Fraction-dependent bounds

The previous section provides a bound that only depends on the types in the population, and not in their relative frequency. However this bound may be too lax if there are many types, or when there is a small fraction of agents with high BPoA.

In this section we provide our main result w.r.t. diverse populations with general behavioral biases.

In the next proposition and theorem, each c_e is (λ, μ)-smooth, each $\hat{c}^i$ is (λ^i, μ^i) smooth, and c_e is $(\hat{\lambda}^i, \hat{\mu}^i)$-biased-smooth w.r.t. $\hat{c}^i$. Recall that $\frac{n_i}{n}$ is the fraction of agents of type i.

PROPOSITION 4.5. *Suppose that game with biased costs* $\mathcal{G} = \left\langle G, \hat{G} \right\rangle$ *is a DSPG with* convex cost functions. *Let* $\mathbf{s}^*$ *be an equilibrium of* $\mathcal{G}$. *Then for any type* i, $SC_i(G, \mathbf{s}^*) \leq \frac{n_i}{n} \frac{\lambda \lambda^i \hat{\lambda}^i}{(1-2\mu)(1-\mu^i)(1-\hat{\mu}^i)} SC(G, OPT(G^i))$.

If the game is also symmetric, then $OPT(G^i) OPT(G)$ for all i. As an immediate corollary we get a bound on the BPoA of $\mathcal{G}$.

THEOREM 4.6. *Suppose that game with biased costs* $\mathcal{G} = \left\langle G, \hat{G} \right\rangle$ *is a* symmetric *DSPG with* convex cost functions. *Then* $\mathit{BPoA}(\mathcal{G}) \leq \sum_i \frac{n_i}{n} \frac{\lambda \lambda^i \hat{\lambda}^i}{(1-2\mu)(1-\mu^i)(1-\hat{\mu}^i)} = O(\sum_i \frac{n_i}{n} \frac{\hat{\lambda}^i}{1-\hat{\mu}^i})$.

In other words, the BPoA of a diverse population on a DSP network is a weighted average of their worst-case PoAs (multiplied by a constant that does not depend on the network or the population structure).

Example 4.7. Consider a bias that maps affine costs to affine costs (e.g., tax sensitivity). We have that $\mu = \mu^i = \frac{1}{4}$, thus

$$\text{BPoA}(\mathcal{G}) \leq \frac{8}{3} \sum_i \frac{n_i}{n} \frac{\hat{\lambda}^i}{(1-\hat{\mu}^i)}.$$

Similarly, for quadratic costs $\mu = \mu^i \cong 0.385$, thus $\text{BPoA}(\mathcal{G}) \leq 7.05 \sum_i \frac{n_i}{n} \frac{\hat{\lambda}^i}{(1-\hat{\mu}^i)}$.

Proof of the main result. Let G^* be a game with biased costs derived from G, where the biased costs are defined as $c^*(x) = c(x) + c'(x)x$.

LEMMA 4.8. *For every $x > 0$, and every type i, it holds that:*

(1) $\frac{\hat{c}^i(x)}{c(x)} \geq (1 - \hat{\mu}^i)$;
(2) $\frac{\hat{c}^i(x)}{c^*(x)} \leq \hat{\lambda}^i$.

PROOF. Define quantity $x' = 0$, so that:

$$x(c(x) - \hat{c}^i(x)) = c(x)x + \hat{c}^i(x)(x' - x) \leq \hat{\lambda}^i c(x')x' + \hat{\mu}^i c(x)x = \hat{\mu}^i c(x)x.$$

Hence $c(x) - \hat{c}^i(x) \leq \hat{\mu}^i c(x)$, and inequality (1) follows.

For (2), consider $x' = x + \Delta$. Then:

$$\begin{aligned}
&c(x)x + \hat{c}^i(x)(x' - x) \leq \hat{\lambda}^i c(x')x' + \hat{\mu}^i c(x)x && \Rightarrow \\
&\hat{c}^i(x)\Delta \leq \hat{\lambda}^i(c(x+\Delta)(x+\Delta)) + (\hat{\mu}^i - 1)c(x)x \leq \hat{\lambda}^i(c(x+\Delta)(x+\Delta)) && \Rightarrow \text{ (since } \hat{\mu}^i \leq 1) \\
&\hat{c}(x) \leq \hat{\lambda}^i(\frac{c(x+\Delta)}{\Delta}x + c(x+\Delta))
\end{aligned}$$

Then we get: $c(x)x + \hat{c}^i(x)(x' - x) \leq \lim_{\Delta \to 0} \hat{\lambda}^i(\frac{\partial c(x)}{\partial x}x + c(x)) = \hat{\lambda}^i c^*(x)$. □

LEMMA 4.9. *Suppose that c is (λ, μ)-smooth, and convex. Then $\forall x\ c^*(x) \leq \frac{\lambda}{1-2\mu}c(x)$.*

PROOF. For $x' = 2x$, we have from smoothness $c(2x)x \leq \lambda c(x)x + \mu c(2x)2x$, and thus $(1 - 2\mu)c(2x) \leq \lambda c(x)$. Therefore, $c^*(x) = c(x) + xc'(x) \leq c(2x) \leq \frac{\lambda}{1-2\mu}c(x)$, where the first inequality is due to convexity. □

This not true for non-convex functions, even if $c(x)x$ is convex. For example, for $c(x) = \sqrt{x}$ and values close to 0, the derivative (and thus $c^*(x)$) goes to infinity. For linear functions the bound is tight. Consider for example $c(x) = x$ (for which $\lambda = 1, \mu = \frac{1}{4}$), then $c^*(x) = c(2x) = 2x = \frac{1}{1-2\cdot\frac{1}{4}}x = \frac{\lambda}{1-2\mu}x = \frac{\lambda}{1-2\mu}c(x)$.

PROOF OF PROP. 4.5. In equilibrium $\mathbf{s}^*$ of $\hat{G}$, all type i agents suffer the same biased cost $F = \hat{c}^i(f, \mathbf{s}^*)$ on all $f \in S_i^*$. In $\mathbf{s}^i$ (the equilibrium of $\hat{G}^i$) the cost for all agents is $F^i = \hat{c}^i(f, \mathbf{s}^i)$ where f is used path (in S^i). Take an arbitrary path $f' \in S_i^* \cap S^i$, then by Lemma 4.1, $s_f^* \leq s_f^i$ and thus $F \leq F^i$.

We denote the optimal state by $OPT(G^i) = (s_f^o)_{f \in A}$. Summing over the cost of all agents in $\mathbf{s}^i$, and using smoothness of $\hat{G}^i$:[6]

$$\begin{aligned}
nF \geq nF^i = (\sum_{f \in A} s_f^i)F^i = \sum_{f \in S^i} s_f^i \hat{c}^i(f, \mathbf{s}^i) = SC(\hat{G}^i, \mathbf{s}^i) \\
\leq \frac{\lambda^i}{1 - \mu^i} SC(\hat{G}^i, OPT(G^i)).
\end{aligned}$$

[6]The reason we cannot use $OPT(G)$ is that $\mathbf{s}^i$ is a biased equilibrium of G^i, and is not a valid profile in the game G (unless we assume symmetry). We thank Eva Tardos for pointing this out.

By definition of the social cost,

$$nF \le \frac{\lambda^i}{1-\mu^i}\sum_e s_e^o \hat{c}_e^i(s_e^o) \le \frac{\lambda^i}{1-\mu^i}\sum_e s_e^o \hat{\lambda}^i \hat{c}_e^*(s_e^o) \qquad \text{(from Lemma 4.8)}$$

$$\le \frac{\lambda^i}{1-\mu^i}\sum_e s_e^o \hat{\lambda}^i \frac{\lambda}{1-2\mu} c_e(s_e^o) = \frac{\lambda\lambda^i\hat{\lambda}^i}{(1-2\mu)(1-\mu^i)} SC(G, OPT(G^i)) \qquad \text{(from Lemma 4.9)}$$

On the other hand, we lower bound the total experienced cost $n_i F$ in terms of the actual cost of $\mathbf{s}^*$. For any $f \in S_i^*$, by Lemma 4.8

$$F = \hat{c}^i(f, \mathbf{s}^*) = \sum_{e\in F} \hat{c}_e^i(s_e^*) \ge \sum_{e\in F}(1-\hat{\mu}^i)c_e(s_e^*) = (1-\hat{\mu}^i)c(f, \mathbf{s}^*).$$

Combining the above bounds,

$$\begin{aligned} SC_i(G, \mathbf{s}^*) &= \sum_{f\in A} s_{i,f}^* c(f, \mathbf{s}^*) \le \sum_{f\in A} s_{i,f}^* \frac{F}{1-\hat{\mu}^i} \\ &= \frac{F}{1-\hat{\mu}^i}\sum_{f\in A} s_{i,f}^* = \frac{F}{1-\hat{\mu}^i} n_i = \frac{n_i}{n(1-\hat{\mu}^i)} nF \\ &\le \frac{n_i}{n}\frac{\lambda\lambda^i\hat{\lambda}^i}{(1-2\mu)(1-\mu^i)(1-\hat{\mu}^i)} SC(G, OPT(G^i)), \end{aligned}$$

as required. □

4.4. Necessity of restricted structure

Lemma 4.1 requires that the network is series-parallel. Without this, even a small fraction of biased agents can increase the cost for all others by an arbitrary factor.

PROPOSITION 4.10. *Consider any graph (V, E) that is not directed series-parallel. Then for any $\epsilon > 0$ and any $M' > 0$, there are symmetric games $G, \hat{G}$ with equilibria $\mathbf{s}^*, \hat{\mathbf{s}}^*$ respectively, such that:*

- *Only a fraction ϵ of the agents in $\hat{G}$ are biased.*
- *There are edges used in $\mathbf{s}^*$ such that $\hat{s}_e^* > s_e^*$.*
- *For the unbiased type i, $SC_i(G, \hat{\mathbf{s}}^*) \ge M' \cdot SC_i(G, \mathbf{s}^*)$.*

PROOF. It is known that 2-terminal network is directed series-parallel, if and only if it contains the Braess graph (the graph that appears in Fig. 3) as a topological minor [Milchtaich 2006]. Thus it is sufficient to construct a bad example on the Braess graph to prove the proposition.

Indeed, consider the network in Figure 3, where the value M is some large number that will be set later on. Suppose the total mass is one unit, and that all agents have need to flow from u to v. In the game G, costs are as in Figure 3 (a). Thus in the equilibrium $\mathbf{s}^*$, the agents will split evenly among the paths u-a-v and u-b-v, so that each agent incurs the same cost of $(2 \cdot 1/2)^M = 1^M = 1$, and thus $SC_1(\mathbf{s}^*, G) = 1 - \epsilon$. Now, denote the unbiased agents as type 1, and suppose that some fraction $\epsilon > 0$ of agents is of a type 2, with biased costs $\hat{c}^2$. See Figure 3 (b).

In the equilibrium $\hat{\mathbf{s}}^*$, all type 2 agents will select the path u-a-b-v, whereas all type 1 agents will split evenly. Thus there will be $\hat{s}_e^* = \frac{1}{2} + \frac{\epsilon}{2} > s_e^*$ agents on each edge s-u and b-v. Note that this already shows that Lemma 4.1 does not apply for our game.

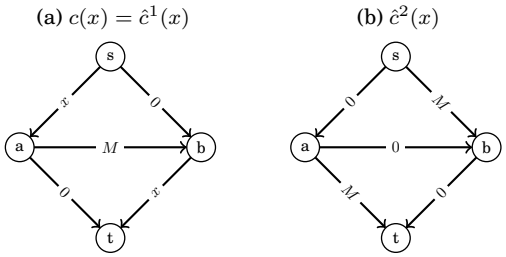

Figure 3: All agents need to select an u-v path. Some fraction of type 2 agents adopt the modified cost function $\hat{c}^2$, where $M \gg 1$ is some large constant.

We now compute the cost for the type 1 agents in $\hat{\mathbf{s}}^*$. Set $M > \frac{M'}{\epsilon}$. For every such agent taking the path $s - a - t$, the cost is

$$c_f(\hat{\mathbf{s}}^*) = c_{s-a}(\hat{s}^*_{s-a}) = c_{s-a}(2(\frac{1}{2} + \frac{\epsilon}{2}))^M = (1+\epsilon)^M > M'.$$

Thus the total cost for type 1 agents is $SC_1(G, \hat{\mathbf{s}}^*) > (1-\epsilon)M' = M' SC_1(G, \mathbf{s}^*)$, as required. □

Since the cost function used in this proof are not smooth, this does not violate the BPoA bound in Theorem 4.6. See Proposition A.1 in the appendix for another example with smooth costs and unbounded BPoA.

5. BIASED SMOOTHNESS FOR TAX-SENSITIVE AGENTS

In this section we consider the role of behavioral bias in the context of tax payments. In particular, suppose that the center imposes the optimal tax of $c'_e(x)x$ on every edge e. For an agent with tax sensitivity $\beta \geq 0$, we have $\hat{c}^\beta(x) = c(x) + \beta c'(x)x$, and the monetary part of the cost is adjusted by a factor of β. We will assume that all agents have the same bias, and provide bounds on $\text{BPoA}(G, \beta) \equiv \text{BPoA}(G, \hat{G}^\beta)$. That is, to bound the efficiency loss in games where all agents have a bias β. When agents differ in their bias the results can be combined with the analysis of bounds under diverse populations.

We consider four classes of games by generality of the cost functions: all standard cost functions, convex functions, polynomial functions with positive coefficients, and affine functions. For example, G_{affine} stands for an arbitrary game with affine costs. Unlike in the previous section, all of our BPoA bounds hold regardless of the network topology. The results are summarized in Tables I and II.

We consider a class of cost functions that are (λ, μ)-smooth, and recall that w.l.o.g. $\lambda = 1$ [Correa *et al.* 2008]. Clearly for $\beta = 0$, the players ignore the tax, and we get $(\hat{\lambda}, \hat{\mu}) = (\lambda, \mu) = (1, \mu)$, $\text{BPoA}(G, \beta = 0) = \frac{1}{1-\mu}$; for $\beta = 1$ we get $(\hat{\lambda}, \hat{\mu}) = (1, 0)$, $\text{BPoA}(G, \beta = 1) = 1$ as explained in Section 3.

PROPOSITION 5.1. *Consider a game where all cost functions are (λ, μ)-smooth. For tax sensitivity $\beta \in [0, 1]$, we have $\hat{\lambda} \leq \beta + (1-\beta)\lambda, \hat{\mu} \leq (1-\beta)\mu$. Thus BPoA$(G_{\lambda,\mu}, \beta) \leq \frac{\beta + (1-\beta)\lambda}{1-(1-\beta)\mu}$.*

PROOF. Let $c = c_e, e \in E$, and $x, x' \geq 0$.

$$\begin{aligned}
c(x)x+\hat{c}^{\beta}(x)(x'-x) &= c(x)x + [c(x) + \beta c'(x)x](x'-x) \\
&= c(x)x + [(1-\beta)c(x) + \beta c(x) + \beta c'(x)x](x'-x) \\
&= (1-\beta)c(x)x + \beta c(x)x + (1-\beta)c(x)(x'-x) + \beta[(c(x)+c'(x)x)(x'-x)] \\
&= (1-\beta)[c(x)x + c(x)(x'-x)] + \beta[c(x)x + c^*(x)(x'-x)] \\
&= (1-\beta)c(x)x' + \beta[c(x)x + c^*(x)(x'-x)] \\
&\leq (1-\beta)(\lambda c(x')x' + \mu c(x)x) + \beta[c(x)x + c^*(x)(x'-x)] && \text{(smoothness of } c(x)\text{)} \\
&\leq (1-\beta)(\lambda c(x')x' + \mu c(x)x) + \beta(c(x')x') && \text{(from Lemma 3.6)} \\
&= \beta + (1-\beta)\lambda c(x')x' + (1-\beta)\mu c(x)x = c(x')x' + (1-\beta)\mu c(x)x.
\end{aligned}$$

The BPoA bound follows immediately from Cor. 3.5. □

We note that in nonatomic games $\lambda = 1$, and hence $\beta + (1-\beta)\lambda = 1$ as well.

PROPOSITION 5.2. *Consider an arbitrary game G. For tax sensitivity $\beta \geq 1$, we have $\hat{\lambda} = \beta, \hat{\mu} = 0$, and BPoA$(G, \beta) \leq \beta$.*

PROOF. Denote $h(x) = xc(x)$, and recall that $h'(x) = c(x) + c'(x)x$.

$$\begin{aligned}
c(x)x+\hat{c}^{\beta}(x)(x'-x) &= c(x)x + [c(x) + \beta c'(x)x](x'-x) \\
&= c(x)x + \beta[c(x) + c'(x)x](x'-x) + (1-\beta)c(x)(x'-x) \\
&= c(x)x + \beta h'(x)(x'-x) + (1-\beta)c(x)x' - c(x)x + \beta c(x)x \\
&\leq c(x)x + \beta h'(x)(x'-x) - c(x)x + \beta c(x)x && \text{(since } \beta \geq 1\text{)} \\
&= \beta[h(x) + h'(x)(x'-x)] \leq \beta h(x') = \beta c(x')x'. \quad \square && \text{(by convexity of } h\text{)}
\end{aligned}$$

PROPOSITION 5.3. *Consider a game $G_{conv,\mu}$ where of cost functions that are both convex $(1,\mu)$-smooth (recall that $\lambda = 1$ w.l.o.g.). For $\beta \geq 1$, we have $\hat{\lambda} = 1 + (\beta-1)\mu, \hat{\mu} = \beta - 1$, and thus BPoA$(G_{conv,\mu}, \beta) \leq 1 + (\beta-1)\mu$.*

PROOF.

$$\begin{aligned}
c(x)x+\hat{c}^{\beta}(x)(x'-x) &= c(x)x + [c(x) + \beta c'(x)x](x'-x) = c(x)x + [c^*(x) + (\beta-1)c'(x)x](x'-x) \\
&= c(x)x + c^*(x)(x'-x) + (\beta-1)c'(x)x(x'-x) \leq c(x')x' + (\beta-1)c'(x)x(x'-x).
\end{aligned}$$

Now, since c is convex, $c'(x)(x'-x) \leq c(x') - c(x)$. Thus,

$$\begin{aligned}
c(x)x+\hat{c}^{\beta}(x)(x'-x) &\leq c(x')x' + (\beta-1)(c(x') - c(x))x = c(x')x' - (\beta-1)c(x)x + (\beta-1)[c(x')x] \\
&\leq c(x')x' - (\beta-1)c(x)x + (\beta-1)[\lambda c(x)x + \mu c(x')x'] \\
&= (1 + (\beta-1)\mu)c(x')x' + (\beta-1)(\lambda-1)c(x)x.
\end{aligned}$$

Again, note that in all NRGs $\lambda = 1$ and thus $\hat{\mu} = 0$.

In particular, we get a BPoA bound of $1 + \frac{\beta-1}{4}$ for affine costs, and $1 + \frac{2(\beta-1)}{3\sqrt{3}}$ for quadratic costs.

5.1. Affine costs

We can derive tighter smoothness bounds for more specific classes of cost functions and biases. The results for affine cost functions are summarized in Table II and in Figure 4.

C	$(1,\mu)$-smooth		convex, $(1,\mu)$-smooth
β	$\beta \leq 1$ [Prop. 5.1]	$\beta \geq 1$ [Prop. 5.2]	$\beta \geq 1$ [Prop. 5.3]
$\hat{\lambda}$	1	β	$1+(\beta-1)\mu$
$\hat{\mu}$	$(1-\beta)\mu$	0	0
PoA upper bound	$\frac{1}{1-(1-\beta)\mu}$	β	$1+(\beta-1)\mu$

Table I: Biased-smoothness bounds for tax-sensitivity

C	affine		quadratic	
β	$\beta \leq 1$ [Prop. 5.4]	$\beta \geq 1$ [Prop. 5.5]	$\beta \leq 1$ [Prop. 5.6]	$\beta \geq 1$ [Prop. 5.7]
$\hat{\lambda}$	1	$\frac{(1+\beta)^2}{4\beta}$	1	$1+\frac{2}{3\sqrt{3}}(\beta-1)$
$\hat{\mu}$	$\frac{(\beta+1)^2}{4}-\beta$	0	$\frac{2}{3\sqrt{3}}\gamma^3-\gamma^2+1$	0
PoA	$\frac{1}{(\beta+1)-\frac{1}{4}(\beta+1)^2}$	$\frac{(1+\beta)^2}{4\beta}$	$\frac{1}{\gamma^2-\frac{2}{3\sqrt{3}}\gamma^3}$	$\frac{(1+2\beta)^3}{27\beta^2}$

Table II: Biased-smoothness bounds for tax-sensitivity, affine and quadratic cost functions ($\gamma = \sqrt{1+2\beta}$)

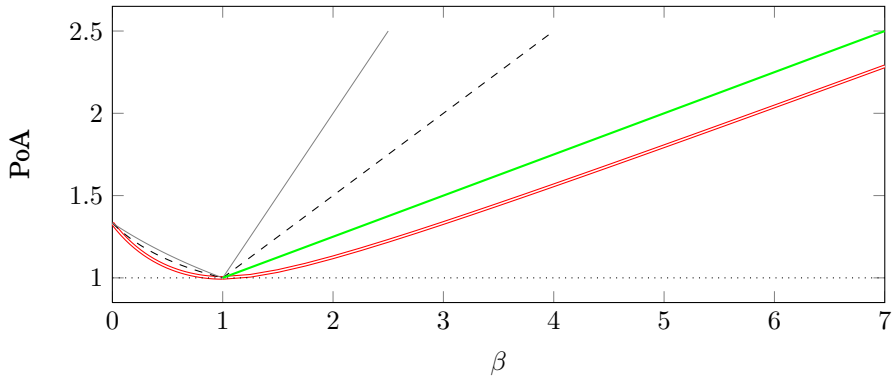

Figure 4: The double red lines are the tight bounds on the PoA in NRGs with affine costs, for agents with tax-sensitivity of $\beta \geq 0$ (see Props. 5.4 and 5.5). For reference, we also show the upper bounds derived for arbitrary cost functions (in light gray, from Props. 5.1 and 5.2); using convexity only (thick green line, from Prop. 5.3); and from Meir and Parkes [2015] (dashed). All bounds apply to pessimistic and risk-averse agents (see Section 6).

While the bounds for affine functions follow as a special case from Prop. 5.6 and Prop. 5.7, we lay out the (shorter) proofs for intuition. We need the following inequality, which is often used in the analysis of nonatomic games:

$$X \cdot Y \leq X^2 + \frac{1}{4}Y^2. \tag{6}$$

Recall that for an affine function $c \in C_{\text{aff}}$, $c(x) = ax + b$, and $\hat{c}^\beta(x) = c(x) + \beta x c'(x) = ax + b + \beta ax = (1+\beta)ax + b$.

PROPOSITION 5.4. *For tax sensitivity $\beta \in [0,1]$, affine cost functions are $(1, \frac{(\beta+1)^2}{4} - \beta)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{aff}, \beta \leq 1$, we have BPoA$(G_{aff}, \beta) \leq \frac{4}{4(\beta+1)-(\beta+1)^2}$. This is a tight bound.*

PROOF. We need to show that $c(x)x + c^\beta(x)(x'-x) \leq c(x')x' + \frac{(\beta-1)^2}{4}c(x)x$.
We have:

$$\begin{aligned}
c(x)x+c^\beta(x)(x'-x) &= ax^2+bx+(a(\beta+1)x+b)(x'-x) = -\beta ax^2 + bx' + a(1+\beta)xx' \\
&\leq -\beta ax^2 + bx' + a[(x')^2 + ((1+\beta)x)^2/4] \text{ (by Eq. (6) with } X = x', Y = (1+\beta)x\text{)} \\
&= (ax'+b)x' + (-\beta + \frac{(1+\beta)^2}{4})ax^2 \\
&\leq (ax'+b)x' + (-\beta + \frac{(\beta+1)^2}{4})(ax^2+bx) \qquad \text{(Since } -\beta + \tfrac{(\beta+1)^2}{4} \geq 0 \text{ for } \beta \leq 1\text{)} \\
&= (ax'+b)x' + (\frac{(\beta+1)^2}{4} - \beta)(ax+b)x = c(x')x' + (\frac{(\beta-1)^2}{4} - \beta)c(x)x.
\end{aligned}$$

We get the upper bound on the BPoA from biased-smoothness and Corollary 3.5. The lower bound is attained by the Pigou example $G_P(a,1)$, for $a = \frac{1}{\beta+1}$. This follows from the analysis in Meir and Parkes [2015], as we explain in the next section, but we directly derive the lower bound for completeness in the appendix. □

For $\beta \geq 1$, we have an upper bound from Prop. 5.3 since affine functions are convex. Thus BPoA$(G_{\text{aff}}, \beta) \leq 1 + \frac{1}{4}\beta$. However we can derive a tighter bound directly.

PROPOSITION 5.5. *For tax sensitivity $\beta \geq 1$, affine cost functions are $(\frac{(1+\beta)^2}{4\beta}, 0)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{aff}, \beta \geq 1$, BPoA$(G_{\text{aff}}, \beta) \leq \frac{(1+\beta)^2}{4\beta}$. This bound is tight.*

PROOF.

$$\begin{aligned}
c(x)x+\hat{c}^\beta(x)(x'-x) &= ax^2+bx+a(1+\beta)xx'+bx'-a(1+\beta)x^2-bx \\
&= a(\frac{(1+\beta)x'}{\sqrt{\beta}})(\sqrt{\beta}x) + bx' - \beta ax^2 \\
&\leq a(\frac{(1+\beta)^2}{4\beta}(x')^2 + \beta x^2) + bx' - \beta ax^2 \quad \text{(by Eq. (6) with } X = \sqrt{\beta}x, Y = \tfrac{(1+\beta)x'}{\sqrt{\beta}}\text{)} \\
&= \frac{(1+\beta)^2}{4\beta}(a(x')^2 + bx' \leq \frac{(1+\beta)^2}{4\beta}(ax'+b)x' = \frac{(1+\beta)^2}{4\beta}c(x')x' \quad \text{(Since } \beta \geq 1\text{)}
\end{aligned}$$

The lower bound is attained by the Pigou example $G_P(a,1)$, for $a = \frac{2\beta}{(\beta+1)^2}$. □

Higher order polynomials

PROPOSITION 5.6. *For tax sensitivity $\beta \geq 1$, polynomial cost functions of degree d are $\left(1, d\left(\frac{1+d\beta}{1+d}\right)^{\frac{d+1}{d}} - d\beta\right)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{poly(d)}$ and $\beta \leq 1$, we have*

$$\textit{BPoA}\left(G_{\text{poly(d)}}, \beta\right) \leq \frac{1}{1 + d\beta - d\left(\frac{1+d\beta}{1+d}\right)^{\frac{d+1}{d}}}.$$

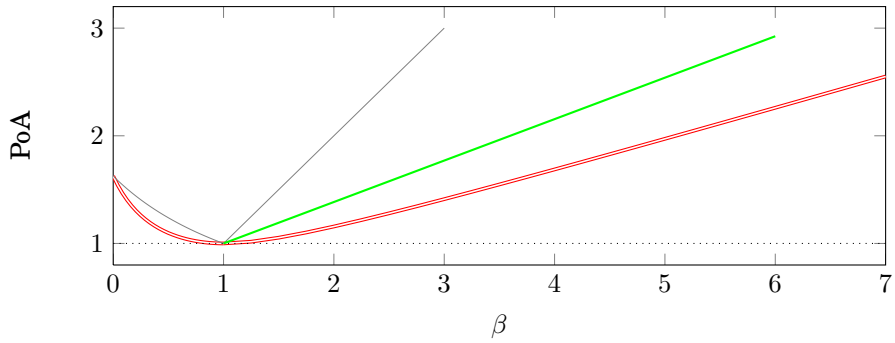

Figure 5: The double red lines are upper and lower bounds on the BPoA in NRGs with quadratic costs, for agents with tax-sensitivity of $\beta \geq 0$ (see Table II). For reference, we show the upper bounds derived for arbitrary cost functions (in light gray, from Props. 5.1 and 5.2); and using convexity only (thick green line, from Prop. 5.3).

This is a tight bound.

Note that for $\beta = 0$, we get the familiar PoA bound of $\left(1 - d(1+d)^{-\frac{d+1}{d}}\right)^{-1} = \Theta\left(\frac{d}{\log(d)}\right)$, whereas for $\beta = 1$ we get that the BPoA is 1.

PROPOSITION 5.7. *For tax sensitivity $\beta \geq 1$, polynomial cost functions of degree d are $(\frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^{d+1}}, 0)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{poly(d)}, \beta \geq 1$, we have $BPoA(G_{\text{poly(d)}}, \beta) \leq \frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^{d+1}}$. This is a tight bound.*

Note that as degree d grows, the bound approaches $\left(1 - \frac{1}{e}\right)\beta$. In contrast, by using Prop. 5.3 we only get a bound of $\mu\beta + O(1)$, which tends to $\beta + O(1)$ as the degree d grows. That is, the bound attainable by only using the convexity property of polynomial costs is greater by a multiplicative factor.

In particular, for both propositions we get tight bounds for the BPoA under quadratic cost functions, which are presented in Table II and in Fig. 5.

6. BIASED SMOOTHNESS FOR THE WCC MODEL

Suppose now that agents are pessimistic, in the sense that they play according to a congestion amount that is larger by a factor of $r > 1$ than the true congestion [Meir and Parkes 2015], that is, $\hat{c}^r(x) = c(rx)$.[7] For affine cost functions, it holds that

$$\hat{c}^r(x) = rax + b = ax + b + (r-1)ax = c(x) + (r-1)c'(x)x,$$

and pessimism with factor $r \geq 1$ coincides with tax-sensitivity of $\beta = r - 1$.

As a result, we can immediately apply the upper bounds for tax-sensitive agents to pessimistic agents with affine cost functions (see Fig. 4).

Lower bounds for pessimistic agents were previously derived in Meir and Parkes [2015], and thus we can also go in the other direction and apply them to the tax-sensitive agents. This correspondence between tax-sensitive agents (for affine costs)

[7] In the model of [Meir and Parkes 2015], agents have uncertainty over the true load, and estimate it to be anywhere in the range $[x/r, rx]$. A risk-averse agent focuses on the upper end of this interval, and hence act as a pessimistic agent.

and risk-averse agents provides an explanation as to why a little uncertainty (or pessimism) initially improves the BPoA, but too much of it hurts: the biased cost increases (due to pessimism) until it reaches the point of optimal taxation, but then goes past it.

Diverse populations. Returning to diverse populations, another result from Meir and Parkes [2015] that we can apply to tax-sensitive agents with affine costs, relates to diverse populations. Consider a game with heterogeneous tax sensitivities $(\beta_i)_{i\leq m}$, and let $\underline{\beta} = \min_i \beta_i$, and $\overline{\beta} = \max_i \beta_i$.

THEOREM 6.1 (DIRECT COROLLARY FROM MEIR AND PARKES [2015]). *For any affine NRG G and tax sensitivities $\boldsymbol{\beta} = (\beta_i)_{i\leq m}$:*

- *If $\beta_i \leq 1$ for all i, then $\textit{BPoA}(G, \boldsymbol{\beta}) \leq \textit{BPoA}(G, \overline{\beta}) = \max_i \textit{BPoA}(G, \beta_i)$.*
- *If $\beta_i \geq 1$ for all i, then $\textit{BPoA}(G, \boldsymbol{\beta}) \leq \textit{BPoA}(G, \underline{\beta}) = \max_i \textit{BPoA}(G, \beta_i)$.*
- *Otherwise, $\textit{BPoA}(G, \hat{G}) \leq \frac{4}{3}\textit{BPoA}(G, \overline{\beta})$.*

We conjecture that a similar technique can be used to prove similar results for any game with a diverse population of tax-sensitive agents, and leave this for future work.

Theorem 4.6 and Theorem 6.1 provide us with two different tools to derive BPoA bounds for diverse populations: the former applies only for DSPGs but takes agents' relative quantities into account, whereas the latter provides a bound in terms of the worst type, but applies for any NRG.

Higher order polynomials. We can apply a similar analysis to other cost functions, although this may lead to different results. Just considering higher order polynomials breaks the correspondence with tax-sensitivity. For example, for quadratic costs:

$$\hat{c}^{\beta}(x) = ax^2 + bx + c + \beta x(2ax + b) = (1 + 2\beta)ax^2 + (1 + \beta)bx + c,$$

whereas $\hat{c}^{r}(x) = a(rx)^2 + b(rx) + c = r^2ax^2 + rbx + c$.

In particular, the BPoA is no longer linear in r. For example, for quadratic cost functions with $r > 2$, the BPoA is between $\frac{r^2}{8}$ and $\frac{r^2}{4}$ (see appendix). More generally for polynomials of degree d, the BPoA is $r^{\Theta(d)}$. In addition, while for some value of $r > 1$ the social cost improves, it does not reach a BPoA of 1, except for affine cost functions.

7. BIASED SMOOTHNESS FOR THE MEAN-VAR MODEL

Following [Nikolova and Stier-Moses 2015], we now consider an arbitrary cost function $c(x)$ and an arbitrary distribution $\epsilon(x)$ s.t. $\mathrm{var}(\epsilon(x)) \leq \kappa c(x)$ for all $x > 0$. Denote $v(x) = \mathrm{var}(\epsilon(x)) \geq 0$, and suppose that $c(x)$ is $(1, \mu)$-smooth for some $\mu < 1$. Define the biased cost as $\hat{c}^{\gamma}(x) = c(x) + \gamma v(x)$.

Nikolova and Stier-Moses bounded the "Price of Risk Aversion" (PRA), which is the ratio between the social welfare in the biased equilibrium and that in the non-biased equilibrium. Their main result is that the PRA is upper bounded by $1 + \kappa\gamma\eta$, where η is a parameter that depends on the network and may be as large as the number of vertices. In particular this leads to a bound of $\text{BPoA}(G_{\mu,\eta}, \gamma, \kappa) \leq (1 + \gamma\kappa\eta)\frac{1}{1-\mu}$ for any $(1, \mu)$-smooth game with network parameter η under risk aversion with parameters γ, κ.

We show the following:

PROPOSITION 7.1. *For $\gamma \geq 0$, $c(x)$ is $(1 + \gamma\kappa, \mu)$-biased-smooth w.r.t. $\hat{c}^{\gamma,\kappa}(x)$.*

PROOF. Suppose first that $x' \leq x$. Then

$$\begin{aligned} xc(c) + \hat{c}^{\gamma,\kappa}(x)(x'-x) &= xc(x) + (c(x) + \gamma v(x))(x'-x) \leq c(x)x + c(x)(x'-x) \\ &= c(x)x' \leq c(x')x' + \mu c(x)x \leq (1+\gamma\kappa)c(x')x' + \mu c(x)x. \end{aligned}$$

(from smoothness of c)

Next, suppose that $x' > x$. Then

$$\begin{aligned} xc(c) + \hat{c}^{\gamma,\kappa}(x)(x'-x) &= xc(x) + (c(x) + \gamma v(x))(x'-x) \leq xc(x) + (c(x) + \gamma\kappa c(x))(x'-x) \\ &= xc(x) + (1+\gamma\kappa)c(x)(x'-x) = -\gamma\kappa xc(x) + (1+\gamma\kappa)c(x)x' \\ &\leq -\gamma\kappa c(x)x + (1+\gamma\kappa)(c(x')x' + \mu c(x)x) \\ &= (1+\gamma\kappa)c(x')x' + \mu c(x)x + \gamma\kappa(\mu-1)c(x)x \qquad \text{(Since } \mu < 1\text{)} \\ &\leq (1+\gamma\kappa)c(x')x' + \mu c(x)x, \end{aligned}$$

as required. □

An immediate corollary is that for any $(1,\mu)$-smooth game G_μ under risk aversion with parameters γ, κ, we have

$$\text{BPoA}(G_{\mu,\eta}, \gamma, \kappa) \leq \frac{1+\gamma\kappa}{1-\mu}.$$

This result improves not just the BPoA bound that follows from Nikolova and Stier-Moses [2015], but also their PRA bound.

Because the unbiased equilibrium is at least at costly as the optimal state, it follows that the PRA is at most the (biased) PoA, and is thus upper bounded by $(1+\gamma\kappa)\frac{1}{1-\mu}$. Crucially, μ is fixed and does not depend on the network structure. This shows that the factor η is not so crucial when considering either the BPoA or the PRA.

We should note however that for some parameter values (namely when η is small and the actual equilibrium of $\hat{G}$ is not much better than the worst-case bound), the bound given by Nikolova and Stier-Moses [2015] is better than ours by a constant factor.

8. DISCUSSION

8.1. Related Work

Various modified cost functions have been considered in the literature of congestion games. A primary example is when taxes (or tolls) are imposed on players using an edge, where the travel time plus the imposed tax can be thought of as a modified cost function. Most related to our work, Bonifaci et al. [2011] design a particular tax mechanism that operates under various restrictions (e.g., taxes may not be too high), and use a smoothness argument to prove BPoA bounds in the resulting game. In a followup work, Jelinek et al. [2014] design optimal taxes for RSGs with restricted taxes, where agents have diverse tax sensitivities. While most work on taxes including the papers above focused on the design of good tax schemes [Cole *et al.* 2003; Fleischer *et al.* 2004], we analyze the affect of a given (standard) taxation scheme on the price of anarchy.

Chen et al. [2011] applied smoothness analysis to provide BPoA bounds for various games (including atomic congestion games) where players are altruistic, i.e., part of their utility is derived from the social welfare. They also bounded the BPoA in games with diverse levels of altruism via direct analysis. We note that in the nonatomic model, altruism amounts to averaging the latency with a number that is constant over all actions, and thus does not affect players' behavior or the BPoA.

[Acemoglu *et al.* 2016] study nonatomic congestion games where some agents are unaware of the existence of certain edges, which is equivalent to having a "wrong"

cost function that assigns infinite costs to some edges. They prove that in DSPGs such ignorance can only lead to a worse equilibrium than under true information, yet worst-case PoA does not exceed the bound of $\frac{\lambda}{1-\mu}$.

Christodoulou et al. [2014] proposed coordination mechanisms for nonatomic congestion games, where the modified cost can be increased arbitrarily. However in contrast with most of the literature on introducing tolls into congestion games, the modified costs are considered "real" rather than "biased," in the sense that the social cost consists of both the latency and the tolls.

Another behavioral bias that has been recently studied in congestion games is risk aversion [Ashlagi *et al.* 2009; Angelidakis *et al.* 2013; Piliouras *et al.* 2013; Nikolova and Stier-Moses 2015; Meir and Parkes 2015]. In some of these models, the behavior of a risk averse agent coincides with that of a rational agent with a particular biased cost function (see Sections 6, 7). We make a detailed comparison in the relevant sections.

Finally, Babaioff et al. [2007] consider congestion games where some of the players are malicious. We can reformulate such a game as a biased game, where malicious behavior corresponds to an *arbitrary bias*.

8.2. Conclusion

We have considered strategic settings in which participants are playing the wrong game, and perceiving utilities in some inaccurate or biased form. Whether these modified utilities, and thus deviations from rational play, come from a cognitive limitation, a behavioral bias, or a different perception in regard to taxes or other payments, it is important to understand how the equilibria of the game are affected.

Biased smoothness is a tool that enables PoA analysis under such modified utilities. Our work is the first to provide PoA bounds on populations with arbitrary biases, diverse across agents, in nonatomic congestion games. The analysis framework reduces the problem to that of analyzing the smoothness parameters of each payoff and behavioral agent type.

Future research directions include deriving PoA bounds under various other types of biases and for more classes of cost functions. A more fundamental question is whether our results on diverse populations with either specific or arbitrary biases can be extended, namely relaxing some of the assumptions in our main results—in particular the restrictions on the graph structure.

Acknowledgments

We thank Evdokia Nikolova, Nicolás Stier-Moses and Eva Tardos for some helpful discussions.

REFERENCES

Heayat Zakaie Aashtiani and Thomas L Magnanti. Equilibria on a congested transportation network. *SIAM Journal on Algebraic Discrete Methods*, 2(3):213–226, 1981.

Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. Informational braess' paradox: The effect of information on traffic congestion. *arXiv preprint arXiv:1601.02039*, 2016.

Haris Angelidakis, Dimitris Fotakis, and Thanasis Lianeas. Stochastic congestion games with risk-averse players. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory*, pages 86–97. Springer, 2013.

Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Two-terminal routing games with unknown active players. *Artificial Intelligence*, 173(15):1441–1455, 2009.

Moshe Babaioff, Robert Kleinberg, and Christos H Papadimitriou. Congestion games with malicious players. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 103–112. ACM, 2007.

M. Beckmann, B. McGuire, and C. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, 1956.

Vincenzo Bonifaci, Mahyar Salek, and Guido Schäfer. Efficiency of restricted tolls in non-atomic network routing games. In *Proceedings of the 4th International Symposium on Algorithmic Game Theory*, pages 302–313, 2011.

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, and Evi Papaioannou. The impact of altruism on the efficiency of atomic congestion games. In *TGC*, pages 172–188. Springer, 2010.

Po-An Chen, Bart De Keijzer, David Kempe, and Guido Schäfer. The robust price of anarchy of altruistic games. In *Proceedings of the 7th Workshop on Internet and Network Economics*, pages 383–390. Springer, 2011.

Giorgos Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 69(3):619–640, 2014.

Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 521–530. ACM, 2003.

José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.

Giovanna Devetag and Massimo Warglien. Playing the wrong game: An experimental analysis of relational complexity and strategic misrepresentation. *Games and Economic Behavior*, 62(2):364 – 382, 2008.

Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 277–285. IEEE, 2004.

Dimitris Fotakis, George Karakostas, and Stavros G Kolliopoulos. On the existence of optimal taxes for network congestion games with heterogeneous users. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory*, pages 162–173. Springer, 2010.

Marguerite Frank. The braess paradox. *Mathematical Programming*, 20(1):283–302, 1981.

Tomas Jelinek, Marcus Klaas, and Guido Schäfer. Computing optimal tolls with arc restrictions and heterogeneous players. In *STACS*, pages 433–444, 2014.

George Karakostas and Stavros G Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 268–276, 2004.

Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th annual conference on Theoretical aspects of computer science*, pages 404–413. Springer-Verlag, 1999.

Reshef Meir and David Parkes. Congestion games with distance-based strict uncertainty. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2015.

Igal Milchtaich. Generic uniqueness of equilibrium in large crowding games. *Mathematics of Operations Research*, 25(3):349–364, 2000.

Igal Milchtaich. Network topology and the efficiency of equilibrium. *Games and Economic Behavior*, 57:321–346, 2006.

Evdokia Nikolova and Nicolas E Stier-Moses. The burden of risk aversion in mean-risk selfish routing. In *Proceedings of the Sixteenth ACM Conference on Economics and*

Computation, pages 489–506. ACM, 2015.

Georgios Piliouras, Evdokia Nikolova, and Jeff S Shamma. Risk sensitivity of price of anarchy under uncertainty. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 715–732. ACM, 2013.

Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.

Tim Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 513–522. ACM, 2009.

William H Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132(1):367–382, 2007.

David Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.

John Glen Wardrop. Road paper. some theoretical aspects of road traffic research. In *ICE Proceedings: Engineering Divisions*, volume 1, pages 325–362, 1952.

Hai Yang and Xiaoning Zhang. Existence of anonymous link tolls for system optimum on networks with mixed equilibrium behaviors. *Transportation Research Part B: Methodological*, 42(2):99–112, 2008.

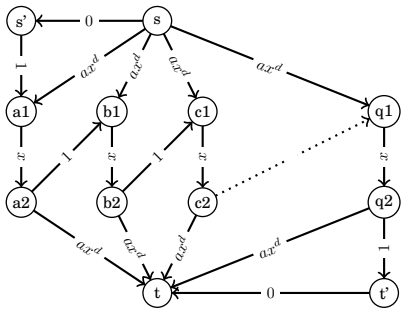

Figure 6: The game G with its base cost functions. There are q disjoint paths of length 3 from s to t.

A. POA LOWER BOUNDS ON NETWORKS

PROPOSITION A.1. *For any $\epsilon \in (0, 0.5)$ and $M \in \mathbb{N}$, there is an NRG (not DSPG) with two types of pessimistic agents, s.t. $n_1 = 1-\epsilon, r_1 = 1; n_2 = \epsilon, r_2 = 8$, and PoA$(G, \hat{G}) > M$.*

PROOF. Consider the network in Figure 6, where the number of parallel paths q is set s.t. $q > \frac{2M}{\epsilon}$, $d > \max\{\log q + 1, 10\}$, and $a = (q/2)^d$. The optimal flow is when agents evenly split among the q paths: then the cost to every agent is

$$2ax^d + x = 2(q/2)^d(1/q)^d + (1/q) = \frac{1}{2^{d-1}} + (1/q) < \frac{2}{q} < \frac{\epsilon}{M}.$$

This is also the social cost. If all agents were of type 1, then the optimal flow would also be the unique equilibrium.

Now consider the behavior of the type 2 agents. The modified cost on the source/target edges becomes $\hat{c}^2(x) = a(rx)^d = a8^d x^d$. The type 1 agents (of mass at least 0.5 still split evenly, thus $x > 1/2q$, and

$$\hat{c}^2(x) > a\frac{8^d}{2^d q^d} = \frac{q^d}{2^d}\frac{8^d}{2^d q^d} = \frac{8^d}{4^d} = 2^d > 2^{\log q + 1} = 2q.$$

So for the type 2 agents, the source/target edges (e.g s-a1) become so expensive, they all must choose the path s-s'-a1-a2-b1-b2-...-q1-q2-t'-t.

Therefore, in equilibrium $\mathbf{s}^*$ we have on each middle edge $\frac{1-\epsilon}{q} + \epsilon > \epsilon$ agents. The (real) cost to each agent is thus at least ϵ, which means $SC(G, \mathbf{s}^*) \geq \epsilon$.

Joining the above results,

$$\text{BPoA}(G, \hat{G}) = \frac{SC(G, \mathbf{s}^*)}{SC(G, OPT)} > \frac{\epsilon}{\epsilon/M} = M.$$

□

B. SMOOTHNESS PROOFS

Proposition 5.4. *For tax sensitivity $\beta \in [0, 1]$, affine cost functions are $(1, \frac{(\beta+1)^2}{4} - \beta)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{aff}, \beta \leq 1$, we have BPoA$(G_{aff}, \beta) \leq \frac{4}{4(\beta+1)-(\beta+1)^2}$. This is a tight bound.*

PROOF OF TIGHTNESS. Note that for any state $\mathbf{s}$, $SC(G, \mathbf{s}) = (1-s) + as^2$.

Consider the Pigou example $G_P(\frac{1}{1+\beta}, 1)$, i.e., where $c_1(x) = 1, c_2(x) = \frac{1}{1+\beta}x$.

Let $s^* = s_2^*$ be the equilibrium load on edge 2. In equilibrium, the biased cost on both edges is the same, thus $1 = \hat{c}_1^\beta(1-s^*) = \hat{c}_2^\beta(s^*) = a(1+\beta)s^* = s^*$. This means $SC(\mathbf{s}^*) = 1 - 1 + a = a = \frac{1}{1+\beta}$. Now consider the state $s' = \frac{1}{2a} = \frac{\beta+1}{2}$. We have that

$$\begin{aligned} SC(OPT) \leq SC(\mathbf{s}') &= 1 - s' + a(s')^2 = 1 - \frac{\beta+1}{2} + \frac{1}{2}s' \\ &= 1 - \frac{\beta+1}{2} + \frac{\beta+1}{4} = 1 - \frac{\beta+1}{4} \qquad \Rightarrow \\ \text{BPoA}(G_P, \hat{G}^\beta) \geq \frac{SC(\mathbf{s}^*)}{SC(\mathbf{s}')} &= \frac{1}{(1+\beta)(1-\frac{\beta+1}{4})} \\ = \frac{1}{\beta+1-\frac{(\beta+1)^2}{4}} &= \frac{4}{4(\beta+1)-(\beta+1)^2}. \end{aligned}$$

□

Proposition 5.5. *For tax sensitivity $\beta \geq 1$, affine cost functions are $(\frac{(1+\beta)^2}{4\beta}, 0)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{aff}, \beta \geq 1$, BPoA$(G_{\text{aff}}, \beta) \leq \frac{(1+\beta)^2}{4\beta}$. This bound is tight.*

PROOF OF TIGHTNESS. The lower bound is attained by the Pigou example $G_P(a,1)$, for $a = \frac{2\beta}{(\beta+1)^2}$.

For the proof, let us denote $r = \beta + 1$, thus $a = \frac{2(r-1)}{r^2}$. As in the previous proposition,

$$s^* = \frac{1}{a(\beta+1)} = \frac{1}{ar} = \frac{r}{2(r-1)}.$$

The equilibrium cost is thus:

$$\begin{aligned} SC(\mathbf{s}^*) &= 1 - s^* + a(s^*)^2 = 1 - \frac{r}{2(r-1)} + \frac{2(r-1)}{r^2}\frac{r^2}{4(r-1)^2} \\ &= 1 - \frac{r}{2(r-1)} + \frac{1}{2(r-1)} = \frac{2(r-1)-r+1}{2(r-1)} = \frac{r-1}{2(r-1)} = \frac{1}{2}. \end{aligned}$$

Consider the state $\mathbf{s}'$ where $s' = 1$. Then $SC(\mathbf{s}') = a$. The BPoA is thus,

$$\text{BPoA}(G_P, \hat{G}^\beta) \geq \frac{SC(\mathbf{s}^*)}{SC(\mathbf{s}')} = \frac{1/2}{a} = \frac{(\beta+1)^2}{4\beta},$$

as required. □

LEMMA B.1. *For any $d > 0$, $\frac{1}{(1+d)^{\frac{1}{d}}} - \frac{1}{(1+d)^{\frac{d+1}{d}}} = \frac{d}{(1+d)^{\frac{d+1}{d}}}$.*

PROOF.

$$\begin{aligned} &\frac{1}{(1+d)^{\frac{1}{d}}} - \frac{1}{(1+d)^{\frac{d+1}{d}}} = \frac{1}{(1+d)^{\frac{1}{d}}} - \frac{1}{(1+d)(1+d)^{\frac{1}{d}}} = \frac{1}{(1+d)^{\frac{1}{d}}}\left(1 - \frac{1}{1+d}\right) \\ &= \frac{1}{(1+d)^{\frac{1}{d}}}\frac{d}{1+d} = \frac{d}{1+d}\frac{1}{(1+d)^{\frac{1}{d}}} = \frac{d}{(1+d)^{\frac{d+1}{d}}}. \end{aligned}$$

□

Proposition 5.6. *For tax sensitivity $\beta \in [0,1]$, polynomial cost functions of degree d are $\left(1, d\left(\frac{1+d\beta}{1+d}\right)^{\frac{d+1}{d}} - d\beta\right)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{poly(d)}$ and $\beta \in [0,1]$, we have*

$$BPoA\left(G_{\text{poly(d)}}, \beta\right) \leq \frac{1}{1 + d\beta - d\left(\frac{1+d\beta}{1+d}\right)^{\frac{d+1}{d}}}.$$

This is a tight bound.

PROOF. We use the inequality [Roughgarden and Tardos 2004]

$$X^d Y \leq d(d+1)^{-\frac{d+1}{d}} X^{d+1} + Y^{d+1}. \tag{7}$$

Let c be a polynomial of degree d^*, and $x, x' \geq 0$.

$$\begin{aligned} c(x)x + \hat{c}^\beta(x)(x'-x) &= \sum_{d=0}^{d^*} a_d x^d x + (\sum_{d=0}^{d^*} a_d x^d + \beta \sum_{d=0}^{d^*} d a_d x^{d-1} x)(x'-x) \\ &= \sum_{d=0}^{d^*} (a_d x^d x' + \beta d a_d x^d x' - \beta d a_d x^{d+1}) \\ &= \sum_{d=0}^{d^*} (a_d (1+d\beta) x^d x' - \beta d a_d x^{d+1}) \end{aligned}$$

For every $d \leq d^*$, setting $X = x(d\beta+1)^{1/d}; Y = x'$ and reorganizing terms,

$$\begin{aligned} a_d(1+d\beta)x^d x' = a_d(X^d Y) &\leq a_d(d(d+1)^{-\frac{d+1}{d}} X^{d+1} + Y^{d+1}) \\ &= a_d(d(d+1)^{-\frac{d+1}{d}}(d\beta+1)^{\frac{d+1}{d}} x^d + x'^{d+1}) \\ &= a_d\left(d\left(\frac{d\beta+1}{d+1}\right)^{\frac{d+1}{d}} x^d + x'^{d+1}\right). \end{aligned}$$

Thus summing over all $d \leq d^*$,

$$\begin{aligned}
c(x)x+\hat{c}^{\beta}(x)(x'-x) &= \sum_{d=0}^{d^*}(a_d(1+d\beta)x^d x' - \beta d a_d x^{d+1}) \\
&\leq \sum_{d=0}^{d^*} a_d \left(d\left(\frac{d\beta+1}{d+1}\right)^{\frac{d+1}{d}} x^d + x'^{d+1} \right) - \beta d a_d x^{d+1} \\
&= \sum_{d=0}^{d^*} a_d d \left(\left(\frac{d\beta+1}{d+1}\right)^{\frac{d+1}{d}} - \beta \right) x^d + x'^{d+1} \\
&\leq \sum_{d=0}^{d^*} a_d d^* \left(\left(\frac{d^*\beta+1}{d^*+1}\right)^{\frac{d^*+1}{d^*}} - \beta \right) x^d + x'^{d+1} \\
&= d^* \left(\left(\frac{d^*\beta+1}{d^*+1}\right)^{\frac{d^*+1}{d^*}} - \beta \right) \sum_{d=0}^{d^*} a_d x^d + \sum_{d=0}^{d^*} x'^{d+1} \\
&= d^* \left(\left(\frac{d^*\beta+1}{d^*+1}\right)^{\frac{d^*+1}{d^*}} - \beta \right) c(x)x + c(x')x'.
\end{aligned}$$

For tightness, consider the Pigou example $G_P(\frac{1}{1+d\beta}, d)$, i.e., where $c_1(x) = 1, c_2(x) = \frac{1}{1+d\beta}x^d$.

The biased cost on edge 2 is $\hat{c}_2^{\beta}(x) = c_2(x) + \beta x c_2'(x) = \frac{1+\beta}{1+d\beta}x^{d+1}$.

Let $s^* = s_2^*$ be the equilibrium load on edge 2. We argue that $s^* = 1$. Indeed, $\hat{c}_2^{\beta}(s^*) = \frac{1+\beta}{1+d\beta}1^{d+1} = \frac{1+\beta}{1+d\beta} \leq 1$, thus no agent wants to switch to edge 1.

This means $SC(\mathbf{s}^*) = 1 - s^* + a(s^*)^{d+1} = a = \frac{1}{1+d\beta}$.

Now consider the state $s' = \frac{1}{((d+1)a)^{1/d}} = \left(\frac{1+d\beta}{1+d}\right)^{1/d}$. We have that

$$SC(OPT) \leq SC(\mathbf{s}') = 1 - s' + a(s')^{d+1} = 1 - \left(\frac{1+d\beta}{1+d}\right)^{\frac{1}{d}} + \frac{1}{1+d\beta}\left(\frac{1+d\beta}{1+d}\right)^{\frac{d+1}{d}}$$

$$\begin{aligned}
\frac{SC(\mathbf{s}')}{SC(\mathbf{s}^*)} &= (1+d\beta)\left(1 - \left(\frac{1+d\beta}{1+d}\right)^{\frac{1}{d}} + \frac{1}{1+d\beta}\left(\frac{1+d\beta}{1+d}\right)^{\frac{d+1}{d}}\right) \\
&= 1 + d\beta - (1+d\beta)\left(\frac{1+d\beta}{1+d}\right)^{\frac{1}{d}} + \left(\frac{1+d\beta}{1+d}\right)^{\frac{d+1}{d}} \\
&= 1 + d\beta - (1+d\beta)^{\frac{d+1}{d}}\left(\frac{1}{(1+d)^{\frac{1}{d}}} - \frac{1}{(1+d)^{\frac{d+1}{d}}}\right) \\
&= 1 + d\beta - (1+d\beta)^{\frac{d+1}{d}}\left(\frac{d}{(1+d)^{\frac{d+1}{d}}}\right) \qquad \text{(By Lemma B.1)} \\
&= 1 + d\beta - d\left(\frac{(1+d\beta)}{(1+d)}\right)^{\frac{d+1}{d}}
\end{aligned}$$

Therefore,

$$\text{BPoA}(G_P, \hat{G}^\beta) \geq \frac{SC(\mathbf{s}^*)}{SC(\mathbf{s}')} = \frac{1}{1 + d\beta - d\left(\frac{(1+d\beta)}{(1+d)}\right)^{\frac{d+1}{d}}},$$

as required. □

Proposition 5.7. *For tax sensitivity $\beta \geq 1$, polynomial cost functions of degree d are $(\frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^{d+1}}, 0)$-biased-smooth w.r.t. $\hat{c}^\beta$. Hence for any game $G_{poly(d)}, \beta \geq 1$, we have BPoA$(G_{\text{poly(d)}}, \beta) \leq \frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^{d+1}}$. This is a tight bound.*

PROOF.
Let c be a polynomial of degree d^*, and $x, x' \geq 0$. As in the previous proof,

$$c(x)x + \hat{c}^\beta(x)(x' - x) = \sum_{d=0}^{d^*}(a_d(1 + d\beta)x^d x' - \beta d a_d x^{d+1})$$

For every $d \leq d^*$, setting $X = x\beta(d+1)^{1/d}; Y = \frac{x'(1+d\beta)}{d+1}$ and reorganizing terms,

$$\begin{aligned} a_d(1 + d\beta)x^d x' &= \frac{a_d}{\beta^d}x^d\beta^d(d+1)\frac{x'(1+d\beta)}{d+1} \\ &= \frac{a_d}{\beta^d}(x\beta(d+1)^{1/d})^d\frac{x'(1+d\beta)}{d+1} = \frac{a_d}{\beta^d}X^dY, \end{aligned}$$

By using inequality (7),

$$\begin{aligned} a_d(1 + d\beta)x^d x' &\leq \frac{a_d}{\beta^d}(d(d+1)^{-\frac{d+1}{d}}X^{d+1} + Y^{d+1}) \\ &= \frac{a_d}{\beta^d}(d(d+1)^{-\frac{d+1}{d}}(x\beta(d+1)^{1/d})^{d+1} + (\frac{x'(1+d\beta)}{d+1})^{d+1}) \\ &= a_d\beta d x^{d+1} + \frac{a_d}{\beta^d}(\frac{x'(1+d\beta)}{d+1})^{d+1}) \\ &= a_d\beta d x^{d+1} + a_d\frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^{d+1}}(x')^{d+1}. \end{aligned}$$

Summing over all $d \leq d^*$,

$$\begin{aligned} c(x)x + \hat{c}^\beta(x)(x' - x) &\leq \sum_{d=0}^{d^*} a_d\frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^{d+1}}(x')^{d+1} \\ &\leq \sum_{d=0}^{d^*} a_d\frac{(1+d^*\beta)^{d^*+1}}{\beta^{d^*}(d^*+1)^{d^*+1}}(x')^{d+1} \\ &= \frac{(1+d^*\beta)^{d^*+1}}{\beta^{d^*}(d^*+1)^{d^*+1}}\sum_{d=0}^{d^*} a_d(x')^{d+1} = \frac{(1+d^*\beta)^{d^*+1}}{\beta^{d^*}(d^*+1)^{d^*+1}}c(x')x'. \end{aligned}$$

For tightness, consider the Pigou game $G_P(\frac{\beta^d(d+1)^d}{(1+d\beta)^{d+1}}, d)$, i.e. $a = \frac{\beta^d(d+1)^d}{(1+d\beta)^{d+1}}$. Denote $\gamma = (1+d\beta)^{1/d}$, then $a = (\frac{\beta(d+1)}{\gamma^{(d+1)}})^d$. In the state $\mathbf{s}'$ where $s'_2 = 1$, we have

$$SC(\mathbf{s}') = 1 - s'_2 + a \cdot (s'_2)^d = a.$$

In equilibrium $\mathbf{s}^*$, we have

$$a \cdot (1+\beta d)(s^*_2)^d = \hat{c}_2(s^*_2) = \hat{c}_1(s^*_1) = 1,$$

meaning that

$$s^*_2 = \frac{1}{a^{1/d}(1+\beta d)^{1/d}} = \frac{\gamma^{(d+1)}}{\beta(d+1)} \frac{1}{(1+\beta d)^{1/d}} = \frac{\gamma^d}{\beta(d+1)}.$$

Finally,

$$\begin{aligned} SC(\mathbf{s}^*) &= 1 - s^*_2 + a(s^*_2)^{d+1} = 1 - \frac{\gamma^d}{\beta(d+1)} + a\frac{\gamma^{d(d+1)}}{\beta^{d+1}(d+1)^{d+1}} \\ &= 1 - \frac{\gamma^d}{\beta(d+1)} + (\frac{\beta(d+1)}{\gamma^{(d+1)}})^d \frac{\gamma^{d(d+1)}}{\beta^{d+1}(d+1)^{d+1}} \\ &= 1 - \frac{\gamma^d}{\beta(d+1)} + \frac{1}{\beta(d+1)} = 1 - \frac{\gamma^d - 1}{\beta(d+1)} \\ &= 1 - \frac{d\beta}{\beta(d+1)} = 1 - \frac{d}{d+1} = \frac{1}{d+1}, \end{aligned}$$

which gives us $\text{BPoA}(G_P, \beta) \geq \frac{SC(\mathbf{s}^*)}{SC(\mathbf{s}')} = \frac{1}{a(d+1)} = \frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^d(d+1)} = \frac{(1+d\beta)^{d+1}}{\beta^d(d+1)^{d+1}}$. □

C. QUADRATIC COSTS WITH PESSIMISTIC AGENTS

PROPOSITION C.1. *For $r \in [1, 1.82]$, quadratic cost functions are $(1, 0.842 - 0.457r)$-biased-smooth w.r.t. $\hat{c}^r$. Hence BPoA$(G_{quad}, r) \leq \frac{1}{0.842-0.457r}$.*

Note that for $r = 1$, we get the familiar bound of $\cong 1.626$ (a rounding of $\frac{3\sqrt{3}}{2-3\sqrt{3}}$. For $r = 1.82$ we get that the BPoA is less than 1.01.

PROOF. We write $c(x) = ax^2 + bx + d$ Then

$$\begin{aligned} c(x)x + c^r(x)(x' - x) &= ax^3 + bx^2 + cx + ar^2x^2x' + brxx' + cx' - ar^2x^3 - brx^2 - cx \\ &= ax^3(1-r^2) + bx^2(1-r) + a(rx)^2x' + b(rx)x' \\ &\leq ax^3(1-r^2) + bx^2(1-r) + a[(x')^3 + \frac{2}{3\sqrt{3}}r^3x^3] + b[(x')^2 + \frac{1}{4}r^2x^2] \\ &= c(x')x' + ax^3[1 - r^2\frac{2}{3\sqrt{3}}r^3] + bx^2[1 - r + \frac{1}{4}r^2] \\ &\leq c(x')x' + (ax^3 + bx^2 + cx)[0.842 - 0.457r] \qquad \text{(for } r \in [1, 1.82]\text{)} \\ &= c(x')x' + (0.842 - 0.457r)c(x)x. \end{aligned}$$

□

PROPOSITION C.2. *For $r \geq 2$, quadratic cost functions are $(\frac{r^2}{4}, 0)$-biased-smooth w.r.t. $\hat{c}^r$. Hence BPoA$(G_{quad}, r) \leq \frac{r^2}{4}$. Also, BPoA$(G_{quad}, r) \geq \frac{r^2}{8}$.*

PROOF.

$$
\begin{aligned}
c(x)x+c^r(x)(x'-x) &= ax^3+bx^2+cx+ar^2x^2x'+brxx'+cx'-ar^2x^3-brx^2-cx \\
&= a(1-r^2)x^3+b(1-r)x^2+2ar^2(x^2)(x'/2)+4br(x)(x'/4)+cx' \\
&\le a(1-r^2)x^3+b(1-r)x^2+2ar^2[\frac{2}{3\sqrt{3}}x^3+(x')^3/8]+4br[\frac{1}{4}x^2+(x')^2/16]+cx' \\
&\le a(1-r^2)x^3+b(1-r)x^2+\frac{4}{3\sqrt{3}}ar^2x^3+ar^2(x')^3/4+brx^2+br(x')^2/4+cx' \\
&\le \frac{1}{4}ar^2(x')^3+\frac{1}{4}br(x')^2+cx' \\
&\le \frac{1}{4}ar^2(x')^3+\frac{1}{4}br^2(x')^2+\frac{1}{4}r^2cx' && \text{(for } r>2\text{)} \\
&= \frac{r^2}{4}c(x')x'
\end{aligned}
$$

For tightness, consider the Pigou game $G_P(\frac{4}{r^2},2)$. □